arXiv:astro-ph/0606615v1 25 Jun 2006

Light Curves from an MHD Simulation of a Black Hole Accretion Disk

Jeremy D. Schnittman

Department of Physics, University of Maryland
82 Regents Drive, College Park, MD 20742

schnittm@umd.edu

Julian H. Krolik

Department of Physics and Astronomy, Johns Hopkins University
Baltimore, MD 21218

jhk@pha.jhu.edu

and

John F. Hawley

Astronomy Department, University of Virginia
P.O. Box 400325, Charlottesville, VA 22904-4325

jh8h@virginia.edu

ABSTRACT

We use a relativistic ray-tracing code to calculate the light curves observed from a global general relativistic magneto-hydrodynamic simulation of an accretion flow onto a Schwarzschild black hole. We apply three basic emission models to sample different properties of the time-dependent accretion disk. With one of these models, which assumes thermal blackbody emission and free-free absorption, we can predict qualitative features of the high-frequency power spectrum from stellar-mass black holes in the "Thermal Dominant" state. The simulated power spectrum is characterized by a power law of index $\Gamma \approx 3$ and total rms fractional variance of $\lesssim 2\%$ above 10 Hz. For each emission model, we find that the variability amplitude should increase with increasing inclination angle. On the basis of a newly-developed formalism for quantifying the significance of quasi-periodic oscillations (QPOs) in simulation data, we find that these simulations are able to identify any such features with (rms/mean) amplitudes $\gtrsim 1\%$ near

the orbital frequency at the inner-most stable orbit. Initial results indicate the existence of transient QPO peaks with frequency ratios of nearly 2:3 at a 99.9% confidence limit, but they are not generic features because at any given time they are seen only from certain observer directions. Additionally, we present detailed analysis of the azimuthal structure of the accretion disk and the evolution of density perturbations in the inner disk. These "hot spot" structures appear to be roughly self-similar over a range of disk radii, with a single characteristic size $\delta\phi = 25^\circ$ and $\delta r/r = 0.3$, and typical lifetimes $T_l \approx 0.3T_{\rm orb}$.

Subject headings: black hole physics – accretion disks – X-rays:binaries

1. INTRODUCTION

Our understanding of accretion onto black holes has grown substantially in recent years. Correlated magneto-hydrodynamic (MHD) turbulence, stirred by an underlying magneto-rotational instability, has now been well-established as the fundamental mechanism of angular momentum transfer in accretion disks (Balbus & Hawley 1998). With that achievement in hand, detailed studies of global accretion disk dynamics have been undertaken via large-scale numerical simulations in both the pseudo-Newtonian (Hawley & Krolik 2001, 2002; Machida & Matsumoto 2003; Armitage & Reynolds 2003) and general relativistic (De Villiers & Hawley 2003; De Villiers, Hawley, & Krolik 2003; Gammie, McKinney, & Toth 2003) frameworks.

At the same time, we have also come into possession of a wealth of observational data. High signal-to-noise spectra of active galactic nuclei (AGN; see e.g., the SDSS composite: Vanden Berk, et al. (2001)) as well as of black hole binaries in a variety of spectral states are now available (McClintock & Remillard 2005); so, too, are detailed light curves and power-density spectra, particularly of black hole binaries (McClintock & Remillard 2005; van der Klis 2005), but also for Seyfert 1 galaxies and other AGN (Markowitz & Edelson 2004). Fluorescent Fe Kα profiles have been used to infer detailed diagnostics of the disk surface in its innermost portions (Fabian et al. 1995; Reynolds & Begelman 1997; Done, Madejski, & Zycki 2000; Miller et al. 2002).

One particularly exciting area of research has been the discovery of high-frequency quasi-periodic oscillations (QPOs) in accreting black hole binaries (Strohmayer 2001a,b). In a growing number of sources, these QPOs appear in pairs with integer frequency ratios of 2:3 (Miller et al. 2001; Remillard et al. 2002; Homan et al. 2005). The high frequencies of these QPOs (near the frequency of the inner-most stable circular orbit; ISCO) suggests a strongly

relativistic origin, but there exists a very broad range of theoretical models, none of which at this point appear overwhelmingly convincing (Abramowicz & Kluzniak 2001; Rezzolla et al. 2003; Rebusco 2004; Schnittman 2005; Petri 2005). At the same time, there have not yet been any clear identifications of QPOs in the MHD simulations, much less pairs of QPOs with integer frequency ratios, making it difficult to choose one imperfect theoretical model over another.

Without a physically consistent way of associating light output with numerical simulations of accretion dynamics, we cannot directly compare the predictions of our best theoretical disk models with any spectral or timing observations. Therefore it is the goal of this paper to move toward the objective of linking dynamical simulations to observational diagnostics. We will do so by applying several phenomenological models to detailed simulation data in order to predict the light that would be produced. From these models, we generate images and light curves for accretion disks as they might be seen by distant observers. Statistical analysis of this derived data will lead us to several interesting generalizations about the nature of variability in the light output of accreting black holes. Although the particular models we employ here for the emissivity and opacity of disks are not entirely realistic, the formalism we create can be readily applied to more realistic models in the future.

2. PRESENT WORK

To be specific, in this paper we will combine the results from general relativistic (GR) MHD simulations similar to those of De Villiers, Hawley, & Krolik (2003) with the ray-tracing and radiative transfer calculations described in Schnittman & Bertschinger (2004) and Schnittman & Rezzolla (2006) to produce light curves and power spectra of accreting black holes. This coupling of ray-tracing with the output data from an independent simulation is often referred to as a "post-processor," a technique commonly used in other branches of physics, particularly when attempting to compare a simulation directly with experiment. The ray-tracing post-processor described in this paper is closely related to a number of radiative transfer codes used in modeling laboratory plasma physics experiments (see, e.g. Pollak et al. (1994); Schnittman & Craxton (2000); King et al. (2001)). While astrophysicists in general do not have the luxury of being able to control (or in many cases even determine) the initial conditions of their experiments, simulations and observational data are achieving levels of quality that will soon prove such post-processors invaluable. One successful post-processor application in astrophysics (for MHD simulations no less!) has been the production of synthetic observations of jets from radio galaxies, which are large enough to have detailed features resolved spatially (Smith et al. 1985; Matthews & Scheuer 1990; Tregillis, Jones, &

Ryu 2001).

Even without imaging observations of accretion disks, the results of the post-processor analysis are very useful in understanding the structure and evolution of the MHD disks, as well as giving valuable insights into X-ray timing measurements of accreting black holes, particularly *RXTE* observations of binary black holes. In this paper we focus primarily on emission models that may track the disk's thermal emission. In future work, we plan to develop other models that will perform the analogous job for coronal emission processes.

The outline and major results of this paper are as follows: we begin in Section 3 with a detailed description of the MHD simulations and the relativistic radiative transfer equation. We describe three different emission models (optically thin line, optically thick line, and thermal blackbody), each useful as a different diagnostic of the accretion disk structure. For the thermal emission model, we describe a method to convert from the dimensionless "code units" of the simulation to more useful cgs units that appear in the absorption and emission coefficients in the radiation transport. In Section 4 we apply these methods to a single MHD simulation of an accretion flow around a Schwarzschild black hole and present light curves and power spectra for the three different emission models and a sample of observer inclination angles. We estimate our sensitivity to detecting high-frequency QPOs in the simulated light curves and discuss the significance and robustness of a transient pair of QPOs with a 2:3 frequency ratio.

To gain a deeper understanding of the power spectra and physical mechanism behind the QPOs, in Section 5 we develop a number of new tools to analyze the azimuthal structure of hot spot perturbations in the turbulent magnetized disk. For a range of radii throughout the disk, we find the characteristic hot spot shape to be self-similar with constant extent in azimuth and radius. Furthermore, we find the hot spot lifetimes to be distributed exponentially, much like the simple model of Schnittman (2005), with a typical lifetime proportional to the orbital period (i.e. the dynamical time) at that radius. Finally, in Section 6, we discuss the implications of these results and outline possibilities for more direct comparison with observations.

3. DESCRIPTION OF THE MODEL

3.1. Data from Global GRMHD Simulations

Our basic data come from a single numerical simulation of accretion onto a Schwarzschild black hole that uses the fully general relativistic 3-dimensional MHD code described in De Villiers & Hawley (2003). This code follows the dynamical evolution of magnetized matter in

the ideal MHD approximation assuming that the fluid's internal energy evolves adiabatically except for the creation of entropy in shocks. Magnetic field evolution utilizes the constrained transport algorithm (Evans & Hawley 1988).

The calculation is done in Boyer-Lindquist coordinates, so the grid's inner radial boundary is placed just outside the event horizon ($r = 2.1M$; throughout this paper, unless stated otherwise, we adopt units in which $G = c = 1$). The outer radial boundary is at $120M$. At both radial boundaries, matter can leave the grid, but not enter it. To avoid the (flat-space) coordinate singularity at the polar axis, there are reflecting boundary conditions at polar angles $\theta = 0.045\pi$ and 0.955π. The full range of azimuthal angles is included in the grid, with, of course, periodic boundary conditions linking $\phi = 0$ and 2π. The 2π range in ϕ is divided into 256 equal-sized zones. However, the 192 zones in both θ and r are concentrated where we expect the sharpest gradients: near the equatorial plane and at small radius (see De Villiers, Hawley, & Krolik (2003) for details).

In the initial state, all the matter is contained within a hydrostatic torus centered on the equatorial plane with pressure maximum at $r = 25M$ and inner edge at $r = 15M$. It possesses a purely poloidal and axisymmetric magnetic field whose fieldlines follow isodensity contours; the plasma initially has a mean β of 100. Outside the torus, the density is fixed at 10^{-10} times the density at the pressure maximum. As found in many simulations before (e.g., De Villiers, Hawley, & Krolik (2003); McKinney & Gammie (2004)), within a few orbits after the beginning of the simulation, the magneto-rotational instability grows to nonlinear amplitude and magnetic stresses redistribute angular momentum throughout the matter. The result is a disk of moderate thickness ($h/r \sim 0.1$) and nearly-Keplerian angular momentum distribution accreting onto the central black hole.

Over the entire simulation run time $T_{\rm sim} = 6000M$, at $5M$ intervals in coordinate time, we write out full 3-dimensional snapshots of the major hydrodynamic variables: the density, pressure, and fluid 4-velocity (in practice we do not actually save data from the first $\sim 1700M$, the time when the early transient features relax). In Figure 1 we plot a selection of time- and azimuth-averaged disk variables from the simulation data (solid curves), and compare them with a steady-state Novikov-Thorne disk (dashed curves). The conversion from code units to cgs units follows the proceedure described in Section 3.3, assuming an accretion rate of $0.5\dot{M}_{\rm Edd}$ onto a black hole of mass $10M_\odot$. Additionally, for the steady-state disk, we have added a small torque at the inner-most stable circular orbit (ISCO). Following Agol & Krolik (2000), we have parameterized this torque by the extra efficiency it adds to the disk through increased dissipation just outside the ISCO. Specifically, in Figure 1 we use $\delta\eta = 0.007$ (chosen to best match the disk height from the simulations; see Fig. 1d), while the no-torque efficiency for a Schwarzschild black hole is $\eta = 0.057$.

The density, temperature, pressure, and radial velocity shown in Figure 1 are mass-weighted averages, to provide a focus on the most "disk-like" region of the MHD simulation near the midplane. Outside $r/M \approx 25$, the simulation cannot really be thought of as an accretion disk, since there is actually net mass flow outwards in that region, the result of the initial torus expanding as angular momentum is transferred from the inner to the outer disk. In the inner-most regions ($r/M \lesssim 10$), the simulation diverges from the Novikov-Thorne disk as the gas acquires a larger inward radial velocity. To maintain mass conservation, this requires a significant decrease of the surface density compared to the Novikov-Thorne disk, as is evident from Figures 1a and 1e.

We should note that it may be somewhat coincidental that the agreement between the simulations and analytic model is so good, as they use rather different equations of state. The simulations include only gas pressure, and ignore any loss of internal energy to radiation, while the Novikov-Thorne model is based on a radiation-pressure dominated disk that balances energy generation by means of viscous dissipation with energy loss by radiation flux from the disk surface. In fact, using significantly different values of $(L/L_{\rm Edd})$ would change the scale height of the Novikov-Thorne disk, while only modifying the cgs scale of the density and pressure of the simulation data and leaving the disk thickness unchanged.

3.2. Ray-tracing and Radiative Transport

The ray-tracing method used in this analysis can be broken up into two independent pieces: (i) the calculation of photon geodesic trajectories through the black hole spacetime, and (ii) the integration of the relativistic radiative transfer equation along these geodesic paths. The first step, which is carried out using a Hamiltonian formulation for the equations of motion, is described in detail in Schnittman & Bertschinger (2004) for the general Kerr metric with Boyer-Lindquist coordinates. The only major difference between our current approach and the methods used in that paper is the manner in which the photon positions and momenta are tabulated. In Schnittman & Bertschinger (2004), the photons passed through a disk of finite thickness, crossing each surface of constant coordinate θ exactly once (see Fig. 1 of that paper). While computationally convenient, this approach limited the range of viewer inclination angles (fully edge-on views were impossible as the observer at infinity would be embedded inside the computational grid and photon intersections with the θ surfaces would occur near the observer and not the accretion disk). Furthermore, since each photon path would terminate after passing through all the θ-surfaces once, we could not produce multiple images via strong gravitational lensing.

We have now solved these problems by tabulating the photon paths according to coor-

dinate time t (the independent variable in our Hamiltonian equations of motion) instead of latitude θ. For each photon, we record the position and momentum at events equally spaced in t over a section of the trajectory centered on the point of closest approach (minimum r) to the black hole (this value of r will always be unique for null geodesics). Figure 2 shows a schematic of this approach for a Schwarzschild black hole, but the discussion below can be applied equally well to the general case of a spinning black hole. This method allows for arbitrary viewing orientations and the modeling of multiply lensed images. Instead of recording the coordinate momenta $p_\mu(\mathbf{x})$, at each point along the path we convert to the locally flat frame of a zero-angular momentum observer (ZAMO), denoted by hatted indices: $p_{\hat{\mu}}(\mathbf{x})$. Working in the ZAMO frame simplifies the solution of the radiative transfer equation, described below.

From the tabulated set of photon positions and momenta we obtain a series of null vectors $dx_i^{\hat{\mu}}$ parallel to $p^{\hat{\mu}}$, defined in the ZAMO frames at space-time coordinates x_i^μ. The fluid properties at these points are interpolated from the fixed computational grid that holds the data output from the MHD simulation. Specifically, we record the fluid 4-velocity $u^{\hat{\mu}}$, also defined in the ZAMO frame, and the gas density and pressure as measured in the fluid frame. As described below in Section 3.3, the code units of density and pressure are converted to cgs units and the temperature is derived assuming a radiation pressure-dominated gas.

With all these pieces in place, the problem is reduced to solving the classical radiative transfer equation through a relativistic fluid with (locally) uniform density, temperature, and velocity. We first write the nonrelativistic radiative transfer equation as

$$\frac{dI_\nu}{ds} = j_\nu - \alpha_\nu I_\nu \ , \tag{1}$$

where ds is the differential path length and I_ν, j_ν, and α_ν are respectively the spectral intensity, emissivity, and absorption coefficient of the fluid medium at a frequency ν. The absorption coefficient is related to the opacity κ_ν through the density ρ: $\alpha_\nu = \rho\kappa_\nu$. In this form, true emission and absorption are included, but not electron scattering. The inclusion of scattering adds an integral over $I_\nu(\cos\theta)$ to the rhs of equation (1), greatly complicating its solution.

Defining the optical depth τ_ν through

$$d\tau_\nu \equiv \alpha_\nu ds, \tag{2}$$

the transfer equation can be written as

$$\frac{dI_\nu}{d\tau_\nu} = S_\nu - I_\nu, \tag{3}$$

where the source function is defined as $S_\nu \equiv j_\nu/\alpha_\nu$. Over regions of constant source function S_ν, equation (3) has the simple solution

$$I_\nu(d\tau_\nu) = S_\nu + e^{-d\tau_\nu}[I_\nu(0) - S_\nu]. \tag{4}$$

Both I_ν and S_ν have the same properties under Lorentz transformations, namely I_ν/ν^3 and S_ν/ν^3 are both invariant. Another Lorentz invariant is the optical depth, since the fraction of photons passing through a finite medium is given by $e^{-\tau}$, which is just a number, and thus the same in any reference frame. Following Rybicki & Lightman (1979), we can calculate the absorption coefficient in a relativistic medium by considering a small volume of matter flowing in the $\mathbf{e}_{\hat{x}}$ direction with respect to the lab frame. Since the motion is in the x direction, the fluid volume thickness l in the y-dimension is the same in the lab (unprimed) and fluid (primed) frames. For a photon propagating at angle θ with respect to the fluid velocity in the lab frame, the total optical depth in the y-dimension can be written

$$\tau_\nu = \frac{l\alpha_\nu}{\sin\theta} = \frac{l}{\nu\sin\theta}\nu\alpha_\nu = \text{Lorentz invariant}. \tag{5}$$

Since $\nu\sin\theta$ is proportional to the p_y component of the photon 4-momentum, it must be the same in both frames because the boost is in a perpendicular direction. Thus $\nu\sin\theta$ is another Lorentz invariant, and so is $\nu\alpha_\nu$. From the definition of the source function $j_\nu = \alpha_\nu S_\nu$, we find that j_ν/ν^2 is also Lorentz invariant, or

$$j_\nu = \left(\frac{\nu}{\nu'}\right)^2 j'_\nu. \tag{6}$$

In the ZAMO frame (here it can be thought of as the lab frame), the spatial path length of the photon trajectory is given by $ds^2 = \eta_{\hat{j}\hat{k}}dx^{\hat{j}}dx^{\hat{k}}$. The photon's 4-momentum and the 4-velocity of the fluid in the ZAMO basis give the angles θ and θ', and the absorption α'_ν and emissivity j'_ν are given in the fluid frame by the specific emission model being used. Now we have enough information to solve the radiative transfer equation in a relativistic flow:

$$\frac{dI_\nu}{ds} = \left(\frac{\nu}{\nu'}\right)^2 j'_\nu - \left(\frac{\nu'}{\nu}\right)\alpha'_\nu I_\nu. \tag{7}$$

The special relativistic redshift between the photons in the ZAMO frame and fluid frame is

$$\frac{\nu}{\nu'} = \gamma(1 + \beta\cos\theta'), \tag{8}$$

where $\gamma \equiv 1/\sqrt{1-\beta^2}$ as usual.

The above analysis, while quite useful for special relativistic flows in the locally flat ZAMO basis, ignores all general relativistic effects of curved spacetime around the black

hole. To include these effects, we need only consider the invariant I_ν/ν^3 along the geodesic path of the photons. This is particularly straightforward from a computational point of view, where the spectrum is stored as a discrete array I^j, evaluated at the frequencies ν^j. These frequencies are redshifted from one point on the trajectory to the next due solely to gravitational effects: Let U_i^μ be the 4-velocity of a ZAMO at position x_i^μ. We can define the inner product with the photon 4-momentum as

$$\chi_i \equiv p_{\mu,i} U_i^\mu. \tag{9}$$

Then the array of frequencies is redshifted along the photon path according to

$$\nu_{i+1}^j = \nu_i^j \left(\frac{\chi_{i+1}}{\chi_i} \right). \tag{10}$$

Similarly, the spectral intensity defined at each frequency point scales as

$$I_{i+1}^j = I_i^j \left(\frac{\chi_{i+1}}{\chi_i} \right)^3. \tag{11}$$

While these methods are well-suited for our implementation of the ray-tracing code, it should be noted that purely covariant approaches also exist for solving the radiative transfer equation in curved spacetime (Fuerst & Wu 2004).

3.3. Emission Models

As in Schnittman & Rezzolla (2006), here we consider three different emission models for the disk. The simplest model, optically thin line emission, assumes that the gas is emitting monochromatic radiation isotropically in the fluid rest frame. The emissivity is proportional to rest mass density, and there is no absorption: $j_\nu \propto \rho\delta(\nu - \nu_{\rm em})$ and $\alpha_\nu = 0$. While not entirely realistic (geometrically thin accretion disks are generally optically thick), this model is extremely useful as a diagnostic tool. The observed photons serve as tracers for the disk's surface density, relativistic beaming, and the gravitational redshift of the fluid near the black hole.

The second model considered is one of optically thick line emission and absorption, with both α_ν and j_ν proportional to $\rho\delta(\nu - \nu_{\rm em})$. Just as the optically thin model served as a diagnostic of surface density perturbations, the optically thick model effectively maps out the Doppler and gravitational red-shifts of the disk's photosphere. In this regard, it can be thought of as a model for the fluorescent iron line seen in many active galactic nuclei and galactic black holes (modulo a specified emissivity profile in r). However, unlike the many

popular thin-disk models such as Laor (1991), our optically thick model samples the upper layers of a vertically extended disk and corona, somewhat complicating the simple iron line profiles predicted by a flat disk.

The third emission model combines a Kramers opacity law with its thermodynamic inverse emissivity:

$$j_\nu \propto \rho^2 T^{-1/2} e^{-x} \tag{12}$$

and

$$\alpha_\nu \propto \rho^2 T^{-7/2} \left(\frac{1 - e^{-x}}{x^3} \right), \tag{13}$$

giving the classical blackbody distribution in an optically thick region of constant temperature:

$$I_\nu = j_\nu / \alpha_\nu \propto T^3 x^3 \frac{e^{-x}}{1 - e^{-x}}, \tag{14}$$

where $x \equiv h\nu/kT$. We believe this final model gives the most accurate description of the "Thermal Dominant," or "High-Soft" State for X-ray binaries.

For the thermal emission model, we must first convert the hydrodynamic variables as provided by the MHD simulations in code units to more useful physical (cgs) units of density, temperature and pressure (as described in De Villiers & Hawley (2003), the fluid variables used in the code are the transport velocity V^i, matter density ρ, pressure P, and Lorentz factor W). To do so requires choosing a few specific model parameters. First, we set the overall time and distance scales of the problem by picking a mass for the central black hole. For the Boyer-Lindquist coordinates used in the simulation, this sets

$$dt_{\rm cgs} = 4.9 \times 10^{-5} dt_{\rm code} M_{10} \text{ sec}, \tag{15a}$$

$$dl_{\rm cgs} = 1.45 \times 10^{6} dl_{\rm code} M_{10} \text{ cm}, \tag{15b}$$

where M_{10} is the black hole mass in units of $10 M_\odot$. To convert code rest-mass density to g/cm^3, we need to set an average mass accretion rate $\dot{M}$ in cgs units:

$$\dot{M}_{\rm cgs} = \dot{M}_{\rm Edd} \left(\frac{L}{L_{\rm Edd}} \right) = \frac{(L/L_{\rm Edd})}{\eta} \frac{L_{\rm Edd}(M)}{c^2}. \tag{16}$$

We then use the code-calculated accretion rate $\dot{M}_{\rm code}$ to get

$$\rho_{\rm cgs} = 2.3 \times 10^{-5} \frac{(L/L_{\rm Edd})}{\eta} \rho_{\rm code} \dot{M}_{\rm code}^{-1} M_{10}^{-2} \text{ g/cm}^3, \tag{17}$$

where the Eddington luminosity is $L_{\rm Edd}(M) = 1.3 \times 10^{39} M_{10}$ erg s^{-1}, and the efficiency η gives the bolometric luminosity through $L = \eta \dot{M} c^2$. The Novikov-Thorne model predicts

$\eta = 0.057$ for the Schwarzschild black hole, and the non-zero torque at the ISCO provides an additional efficiency factor of $\delta\eta = 0.007$.

The cgs pressure and temperature of the gas are determined by assuming a radiation-pressure dominated equation of state, which is quite reasonable for the inner regions of a black hole accretion disk with an appreciable fraction of the Eddington luminosity. For such a gas, the isothermal sound speed is $c_s = (p_r/\rho)^{1/2}$, giving

$$p_{\rm cgs} = \rho_{\rm cgs} c_s^2 = p_{\rm code} \left(\frac{\rho_{\rm cgs}}{\rho_{\rm code}} \right) c^2, \tag{18}$$

from which the temperature is easily derived:

$$T_{\rm cgs} = \left(\frac{3 p_{\rm cgs}}{a} \right)^{1/4}, \tag{19}$$

where a is simply the radiation density constant (7.6×10^{-19} erg cm^{-3} K^{-4}).

As mentioned above in Section 3.1, this method of determining the gas temperature is not entirely consistent with the assumptions made in the MHD simulation. Namely, the code assumes a gas pressure-dominated adiabatic equation of state, with no explicit radiation pressure and the gas entropy increased only by shock heating. Yet perhaps this is a case of the "ends justifying the means:" even though we assume a radiation-dominated pressure to derive the temperature from a calculation that actually ignores radiation pressure, the resulting values are actually in close agreement with the analytic predictions of a standard alpha-disk for the same black hole mass and accretion rate. Alternatively, instead of using the mass accretion rate to determine the code-to-cgs conversion factors, we could attempt to match any of the other Novikov-Thorne disk variables, such as surface density, central pressure, or temperature. These all give somewhat different scaling factors for the density, but we find the behavior of the light curves for the thermal emission model are ultimately very similar in all cases.

As a demonstration of how these different models can be used to probe the structure of the accretion disk, we show in Figure 3 a snapshot of the inner disk, as seen by observers at inclinations of $i = 0°$ and $70°$ ($90°$ is edge-on) for the three emission models. For these figures and throughout the majority of this paper, we include only emission from the inner region of the accretion flow with $r < 25M$, where the gas behaves most like a classical accretion disk. In all panels, the observed intensity is represented by a color-coded logarithmic scale normalized to the peak intensity of that frame. In the high-inclination frames, the gas is moving towards the observer on the left-hand side of the disk, i.e. in the positive $\hat{\phi}$ direction.

In Figures 3a and 3b, the optically thin nature of the gas is clear from the Einstein ring seen in the inner-most region of the disk: this ring is formed by the photons that pass

around the black hole and then through the disk and to the observer. With the exception of the emission from the Einstein ring, the intensity clearly falls off with decreasing radius as the surface density decreases inward with radius and the gravitational redshift further reduces the observed emission. In Figure 3b, the special relativistic beaming of photons is clear from the higher intensity seen from the gas moving towards the observer on the left.

In Figures 3c and 3d, we show the optically thick line emission model, demonstrating the nearly uniform, monochromatic emission from the face-on disk. For the face-on view, the only variations in intensity come from the gravitational redshift and the transverse Doppler shift of gas orbiting more rapidly in the inner-most disk. Small local pertubations are also caused by the Doppler shift due to turbulent motion of the photosphere perpendicular to the plane of the disk. As in Figure 3b, the edge-on view in Figure 3d shows the beaming of photons towards the observer, but produced from the upper layers of the optically thick disk and surrounding atmosphere.

Finally, in Figures 3e and 3f, we show the image of the disk using the thermal emission and absorption model. Unlike the optically thin model where emission traces surface density, here the emission is greatest in the inner disk where the temperature is highest. Also, the non-linear dependence of emission and absorption on the hydrodynamic variables produces larger amplitude modulations, clearly seen in the structure of the disk. Because we include emission and absorption only inside of $r = 25M$, the interior of the disk can be seen in Figure 3f as a bright band at the outer edge. While this effect is not physical (in a real accretion disk it would be blocked by the outer portion of the disk), it nonetheless serves as a useful imaging device and we believe it does not significantly affect the qualitative behavior of the light curve. One piece of evidence in support of this claim is the lack of any significant excess power at the orbital frequency of $r = 25M$ in the high-inclination light curves.

Just as the optically thin emission model can be used to track surface density in the disk, the thermal model is able to measure the local density and temperature at different depths of the disk. In this way, the ray-tracing analysis provides us with both *literal* and *figurative* X-ray vision of the disk. The regions of highest temperature (typically closer to the disk midplane) produce photons in the high-energy tail that can pass directly through the disk to the observer ($\alpha_\nu \sim x^{-3}$ for $x \gg 1$). Conversely, since the disk is optically thick to the low-energy Rayleigh-Jeans tail of the thermal spectrum, the low-energy photon emission is a direct measurement of the temperature at the surface of the disk, where the intensity is proportional to $\nu^2 T$.

This 3-dimensional imaging method is demonstrated in Figure 4 for the same disk snapshot shown in Figure 3e, now divided into three different energy bands: 1, 10, and 100 keV. The thermal emission model has an inner disk temperature $T_{\rm in} \sim 2 - 3$ keV, so the

three plots sample the low- and high-energy tails of the spectrum and the thermal peak. As described above, the 1 keV emission shown in Figure 4a probes the surface temperature of the disk, highlighting the inner-most region where the gas is hottest. The peak of the thermal spectrum near $3kT_{\rm in} \approx 10$ keV (Fig. 4b) of course produces the majority of disk emission and thus more closely resembles the integrated emission shown in Figure 3e. Since the disk is actually optically thin to the high-energy photons at 100 keV, they can be used to trace the regions of greatest temperature and over-density near the midplane (Fig. 4c). We find that these regions are distributed nearly uniformly in radius throughout the disk.

4. LIGHT CURVES AND POWER SPECTRA

Merging the ray-tracing methods described above with the time-dependent output from the MHD simulations, we are able to produce simulated light curves of the accretion disk. For each of the three emission models, we plot in Figure 5 the light curves for three different inclination angles. A total of 870 data frames have been sampled, spaced at coordinate time steps of $5M$, corresponding to a total simulation time of just over 200 ms for a $10M_\odot$ black hole, or about fifty orbits at the ISCO. This is a remarkably small amount of "real" time for a simulation that takes hundreds of hours to run on a major supercomputer. Nonetheless, we believe it is still sufficiently long enough to characterize the broad-band features in the power spectrum, and additionally, as we show below, to provide interesting constraints on the possible presence of high-frequency QPOs.

Each of the light curves in Figure 5 has been normalized and, for inclinations $i = 45°$ and $75°$, shifted vertically in order to facilitate comparison. Additionally, we have applied a "pre-whitening" filter to each light curve, removing linear secular trends and fixing $I(t_i) = I(t_f)$ in order to avoid spurious high-frequency noise in the Fourier power spectra. For the optically thin line emission model, shown in Figure 5a, it is clear that the same large-amplitude, low-frequency features appear in phase for all inclinations. These features most closely represent the total mass within $r < 25M$, a variable that is primarily driven by the mass flux in the outer regions of the disk, and of course is independent of inclination angle. The higher-frequency variability is due to regions of overdensity moving through the gravitational potential and also getting beamed towards and away from the observer for higher inclinations. Thus the short-term variability is greater for inclinations of $45°$ and $70°$, which is evident from inspection of Figure 5a.

For the other emission models, the low-frequency correlation between different inclinations is still present, but weaker than in the optically thin case. This is quite reasonable, as the optically thick models are more sensitive to localized perturbations that appear different

to observers at different inclinations. Similarly, there is clearly more fractional variation on short time scales, again due to the increased importance of local perturbations: if a hot spot rises to the top of the disk and then disappears again on a dynamical time scale, the effect will be greater for an optically thick disk than a transparent one. Furthermore, the light curves from different emission models do not appear to be correlated with each other, confirming the presumption that they probe fundamentally different properties of the disk.

Moving beyond simple visual inspection of the light curves, we plot in Figure 6 the power spectral density, in units of $(\mathrm{rms/mean})^2\mathrm{Hz}^{-1}$, for the same light curves plotted in Figure 5. As suggested by the discussion above, we find significantly more high-frequency power in the optically thick and thermal emission models. Within each model, we find more high-frequency power for the systems with higher inclinations. The total power for each light curve is listed in Table 1, as well as the amount of high-frequency ($f > 100$ Hz) power.

In all the plots shown in Figure 6, the power spectra can be described by a steep power-law in frequency. There is some evidence of a flatter slope below ≈ 50 Hz, but this is quite possibly a spurious effect due to the short duration of the simulation, so that we have poor statistical sampling of low-frequency features in the light curves. We have listed in Table 1 the best-fit power-law indices for each power spectrum above 50 Hz, mostly in the range $\Gamma \sim 3 - 4$. While the spectra are somewhat flatter for the optically thick and thermal models, all appear steeper than those seen in *RXTE* spectra, which are more typically around $\Gamma \sim 1$ (McClintock & Remillard 2005) (however it should be noted that the data become increasingly poor above ~ 100 Hz for most observations in the Thermal state). At the same time, the simulations of Armitage & Reynolds (2003), where emission traces the vertically integrated stress, produce power spectra with slopes of $\Gamma \approx 2$.

In all cases the slope appears to be independent of the inclination angle, and a function only of the emission model. For each emission model, the total rms amplitude of the light curve fluctuations increases with inclination, a result that was also found by Armitage & Reynolds (2003). In fact, this may be the most robust measure with which to compare the simulated light curves with observations. For a broad range of possible emission mechanisms, we find that the MHD turbulent disks consistently appear more variable when viewed at high inclinations, a prediction that should be testable with archived *RXTE* data.

As shown in Armitage & Reynolds (2003), the power spectrum from a single radius (i.e. annulus of the disk) should show a clear break frequency at roughly the orbital frequency at that radius. Following that approach, we show in Figure 7 the power spectra from five different radii in the disk. Again we use the thermal emission model and a face-on inclination $i = 0^\circ$ so as not to introduce spurious features from the direct emission from the disk interior (artificially visible due to cutting off absorption outside the given annulus). The radii

sampled in Figure 7 have $r/M = (2-6)$, $(6-10)$, $(10-15)$, $(15-20)$, and $(20-25)$. There is some visible evidence for broken power laws, with lower break frequencies for annuli at larger radius, but again the short simulation time severely limits our confidence in the behavior of the power spectra at low frequencies. More statistically significant is the difference between the power-law slopes at high frequency, where the inner regions of the disk clearly have flatter power spectra. The inner regions also produce much more total power, dominating the overall variance in the integrated flux. These results are summarized in Table 2, which lists the best-fit power-law indices above and below 100 Hz, as well as the integrated power for each annulus.

Even if we were to include emission only from the inner-most regions of the disk, the corresponding power spectra would still be significantly steeper than those seen in the *RXTE* observations. The discrepancy with the the $\Gamma = 2$ result of Armitage & Reynolds (2003) should also be investigated in future work in order to understand the relative importance of the simulation details compared to the post-processor emission models in producing a given power-law slope. In fact, these discrepancies provide one of the strongest arguments demonstrating the need of an accurate radiation transfer post-processor for the MHD simulations. By looking solely at the variation of the code variables, as many previous works have done, one cannot hope to reproduce the light curves and power spectra that are actually observed from the disk.

To agree quantitatively with the data will require either new physics in the MHD simulations or new emission models for the post-processors, presumably with an increased focus on the emission from the corona. Additionally, we know electron scattering plays an important role in the radiation transport for hot accretion disks, and therefore geodesic photon paths are not completely realistic. As yet further motivation for understanding the emission processes that produce these power spectra, recent observations of GRS 1915+105 actually show a much steeper power-law, with $\Gamma = 2.8 - 3.0$ (Belloni et al. 2006). However, those observations were taken during a relatively hard state (as opposed to the soft Thermal state), further complicating direct comparison with our emission models.

4.1. Sensitivity to QPOs

From visual inspection of the power spectra in Figure 6, no clear features (e.g. peaks, power-law breaks, etc.) are immediately evident. However, the short simulation time makes inherent statistical fluctuations unavoidable, which could either hide true features or falsely imitate absent ones. We would therefore like to derive a more quantitative understanding of our sensitivity to timing features in the simulated light curves. From the power spectra

plotted in Figure 6c, we can derive an estimate for the theoretical upper limits on the possible existence of high frequency QPOs in the MHD simulations.

With the high spatial resolution of our ray-tracing post-processor, we can safely assume the power-law behavior of the power spectrum is due to the inherent turbulence of the disk, as opposed to any sort of "detector" or finite sampling statistics. Further assuming that the power in neighboring frequency bins is uncorrelated, we can treat the variations (jitter) in the power spectrum as behaving locally like Poisson variants. In the limit of an infinite length simulation, the frequency bins would become infinitesimally small, and thus could be averaged to get an arbitrarily smooth power spectrum, i.e., the "inherent" turbulent power, which we call $P_\infty(\nu)$, modeled as a simple power-law:

$$P_\infty(\nu) = P_0 \nu^{-\Gamma}. \tag{20}$$

Then the probability distribution function for the power at a given frequency is given by that of a Poisson variant:

$$f(P_\nu) = \frac{1}{P_0 \nu^{-\Gamma}} \exp\left(\frac{-P_\nu}{P_0 \nu^{-\Gamma}}\right), \tag{21}$$

which has mean $\mu = P_0 \nu^{-\Gamma}$ and variance $\sigma^2 = \mu^2$. We have checked this relation empirically for the power spectra in Figure 6 and find that the measured variations in P_ν do indeed behave according to equation (21).

For this distribution, the observed power in a single frequency bin is expected to be $P_\infty(\nu)$, plus or minus 100% (hence the critical importance of frequency binning). To identify a potential QPO, we would integrate over some range in frequency $\nu - \Delta\nu$ to $\nu + \Delta\nu$. For a simulation duration of time $T_{\rm sim}$, the number of frequency bins in this range is

$$N_{\rm bin} = 2\Delta\nu T_{\rm sim}. \tag{22}$$

Assuming $P_\infty(\nu)$ is roughly constant over this range (i.e. $\Delta\nu \ll \nu$), the total power P_N in the $N_{\rm bin}$ frequency bins should have a distribution than is approximately normal (central limit theorem for large $N_{\rm bin}$):

$$f(P_N) = \frac{1}{\sqrt{2\pi N_{\rm bin}\sigma^2}} \exp\left[\frac{-(P_N - N_{\rm bin}\mu)^2}{2N_{\rm bin}\sigma^2}\right]. \tag{23}$$

A potential QPO in the power spectrum would appear as a small amount of excess power above the "background noise" $P_\infty(\nu)$. A QPO centered around ν_0 with a FWHM of $2\Delta\nu$ and a Lorentzian profile

$$P_{\rm QPO}(\nu) = \frac{A_{\rm QPO}^2 \pi/\Delta\nu}{1 + \left(\frac{\nu-\nu_0}{\Delta\nu}\right)^2} \tag{24}$$

has total (rms/mean) amplitude of

$$A_{\rm QPO}({\rm rms}\ \%) = \sqrt{\int d\nu P_{\rm QPO}(\nu)}. \tag{25}$$

For our significance estimate, we will consider only the power within $\nu_0 \pm \Delta\nu$, which has an rms departure from the background power-law of $A_{\rm QPO}/\sqrt{2}$. From equation (23), we see that the total measured power

$$P_{\rm tot}(\nu_0, \Delta\nu) = \int_{\nu_0-\Delta\nu}^{\nu_0+\Delta\nu} P_\nu d\nu \tag{26}$$

can be used to give a "1-sigma" confidence limit on the QPO power by

$$P_{\rm tot} = 2\Delta\nu P_\infty(\nu_0) + P_{\rm QPO} < (N_{\rm bin}\mu + \sqrt{N_{\rm bin}\sigma^2})d\nu, \tag{27}$$

where $d\nu = 1/T_{\rm sim}$ is the size of the frequency bin in Hertz. More generally, at the "n-sigma" significance level, we can rule out any QPO with rms amplitude greater than

$$\begin{aligned} A_{\rm QPO}(n_\sigma) &= \left[2(N_{\rm bin}\mu d\nu + n_\sigma\sqrt{N_{\rm bin}\sigma^2}d\nu - 2\Delta\nu P_0\nu_0^{-\Gamma})\right]^{1/2} \\ &= \left[2n_\sigma d\nu\sqrt{N_{\rm bin}\sigma^2}\right]^{1/2} = \left[n_\sigma P_0\nu_0^{-\Gamma}\sqrt{\frac{8\Delta\nu}{T_{\rm sim}}}\right]^{1/2}. \end{aligned} \tag{28}$$

In terms of the oscillator quality factor $Q = \nu_0/FWHM = \nu_0/(2\Delta\nu)$, the amplitude limit can be written

$$A_{\rm QPO}(n_\sigma) = \left(2n_\sigma P_0\nu_0^{-\Gamma+1/2}\right)^{1/2}(T_{\rm sim}Q)^{-1/4}. \tag{29}$$

Plugging in some typical numbers from our simulations (e.g., the power spectra in Fig. 6c), we have $P_0 = 10$, $\Gamma = 3$, and $T_{\rm sim} = 200$ msec. For a typical high-frequency QPO with $\nu_0 = 200$ Hz and $\Delta\nu = 20$ Hz ($Q = 5$; Remillard et al. (2002)), we should thus be able to rule out the presence of a QPO with rms amplitude 1% at the 3-sigma significance level.

A number of important relationships can be seen from equation (28). First, the dependence on ν_0 shows that due to the steep power-law behavior of the background turbulence, with the current simulations it is very difficult to detect or rule out QPOs at *low* frequencies. For example, a similar QPO with quality factor $Q = 5$ at $\nu_0 = 50$ Hz would require an amplitude of $\sim 6\%$ to be detected at the same significance level. Narrower QPOs (large Q, small $\Delta\nu$) are easier to detect, but not by much, due to the $\Delta\nu^{1/4}$ dependence in equation (28). Similarly, a much longer simulation would also not improve sensitivity that much, due to the same weak dependence on $T_{\rm sim}$.

4.2. Identification of High-frequency QPOs

The discussion above provides a simple way to estimate *a priori* the amplitude of any potential QPO at a given significance limit. Turning the problem around, and without assuming anything about the shape or amplitude of the QPO, one can simply calculate the significance of any power excess found in a localized group of $N_{\rm bin}$ frequency bins. In this approach, we set a fixed value for the quality factor Q, which in turn gives $N_{\rm bin}(\nu)$ and thus $P_N(\nu)$, from which the significance of any region of excess (or deficit) power can be determined by

$$n_\sigma(\nu) = \frac{P_N(\nu) - N_{\rm bin}\mu}{\sqrt{N_{\rm bin}\sigma^2}}. \tag{30}$$

Note that this definition of n_σ allows for negative significance values (i.e. deficit power) and has a mean value of 0. If the distribution of power in individual frequency bins follows an exponential distribution, which is what we find for our power spectra, the distribution of n_σ will have a chi-squared distribution with $2N_{\rm bin}$ degrees of freedom. For large values of $N_{\rm bin}$, this approaches a normal distribution with variance $N_{\rm bin}\sigma^2$.

In Figure 8 we plot a collection of these significance curves for the thermal emission model with a range of inclination angles and orientations, fixing $Q = 10$. For each power spectrum, we determine the power-law background $P_\infty(\nu)$ with a least-squares fit to the simulated data. While not immediately obvious in the original power spectra, a number of interesting features appear in the significance analysis. Most striking are the peaks in Figure 8c that appear close to the orbital frequency at the ISCO, as well as $2/3\nu_\phi$(ISCO) and the first harmonic $2\nu_\phi$(ISCO). The significance of these peaks increases with inclination (compare with the smaller peaks at the same frequencies in Fig. 8b, where $i = 45°$), as predicted by the hot spot model described in Schnittman & Bertschinger (2004).

This result is very intriguing, especially in light of a number of recent observations of black hole binaries that find multiple QPO peaks with a 2:3 frequency ratio (Miller et al. 2001; Remillard et al. 2002; Homan et al. 2005). Thus it is quite important to understand the robustness of such a result, if possible without requiring hundreds of hours more of supercomputer simulations. As shown below in Section 5, the average density pertubation in the disk has a lifetime of less than a third of an orbital period. Therefore two different light curves of the same disk viewed from $\phi \pm 90°$ should be uncorrelated when seen at high inclination, where the emission is dominated by special relativistic beaming from the half of the disk moving towards the observer. Similarly, for an optically thick disk, an observer looking at the "underside" of the disk ($i \to 180° - i$) might also see an uncorrelated light curve. We calculated light curves and power spectra with $\theta = 70°, 110°$ and $\phi = \pm 90°$, and show the significance plots in Figures 8d-f. The light curves from $\phi = \pm 90°$ are indeed

uncorrelated. On the other hand, the top-bottom views appear to be strongly correlated. From this we infer that the largest hot spot perturbations extend vertically through the entire disk, and are thus visible from both sides simultaneously.

As is clear from Figures 8d and 8f, the 2:3 QPO peaks are not always present, while other peaks may appear at other frequencies. Although the transience of this result is somewhat disappointing, it is actually consistent with the data from *RXTE*, which detects the subtle HFQPOs in only a small fraction of all observations (McClintock & Remillard 2005). On the other hand, if the simulations most accurately represent the Thermal state (which is by no means certain, as evidenced by the large discrepancy in theoretical vs. observed power-law slope), one would not expect to see high-frequency QPOs at all, because they are not generally observed in the state. Yet another caveat is that around 230 Hz, for $T_{\rm sim} = 200$ msec, $N_{\rm bin} \approx 5$, so the distribution of n_σ is not quite normal. Whereas a random variable with a Gaussian distribution would have 4σ outliers only 0.006% of the time, the actual distribution for $N_{\rm bin} = 5$ (roughly chi-squared with 10 degrees of freedom) gives 4σ results with probability 0.2% and 5σ outliers 0.03% of the time. Thus is it unlikely, but not impossible, to form a few large ($> 4\sigma$) peaks from purely random fluctuations.

In order to understand better the likelihood of these results, we performed a Monte Carlo calculation of 10^6 simulated power-law power spectra. This is a similar approach to that taken by people searching for QPOs in X-ray observations of Sgr A* and AGN (Benlloch et al. 2001; Belanger et al. 2006). For fixed values of P_0 and Γ, we generate a random value for the power at each frequency according to the exponential distribution in equation (21). Then we analyze this simulated power spectrum with the exact same methods described above, generating a plot of significance versus frequency like the ones shown in Figure 8. Repeating this procedure many times gives the probability of producing a random QPO signal at a given frequency and significance. For $Q = 10$, Figure 9 plots this probability for $n_\sigma = 3$, 4, 5, and 6. The results of the Monte Carlo calculation appear to be very similar to the analytic estimates presented above, confirming the non-normal nature of the distribution. The "steps" or "edges" seen in Figure 9 are simply due to the finite number of frequency bins used in each QPO: for a fixed value of Q, $N_{\rm bin}$ increases with frequency, and every integer step corresponds to a small change in the significance, as can be seen from equation (30).

In the end, we are still not quite sure if the 2:3 ratio QPOs are "real," and if so, whether they represent the same QPO pairs seen in some observations. Analogous to the observations, the only way really to improve our confidence in these results may be to run more simulations and see if the 2:3 peaks are reproducible and if so, with what range of QPO parameters. With more simulated data, we should also be better able to identify the

exact mechanism that is causing the QPOs. At this point, it seems somewhat unlikely that they are caused by a simple hot spot model as in Schnittman (2005), which requires much longer lifetimes for the quality factors seen here. To understand the robustness of these QPO pairs, in addition to longer simulation times, we must also explore the effects of black hole spin, one of the major ingredients in most HFQPO models. This effort, too, will require significant computational resources, but should be more productive than simply running the current Schwarzschild simulation longer [see eqn. (29)].

5. AZIMUTHAL STRUCTURE OF ACCRETION DISK

To gain a deeper understanding of the source of the light curve fluctuations, we must first characterize the behavior of the hydrodynamic fluctuations in the accretion disk. For all three emission models, and particularly the optically thin model, the variations in brightness are closely related to the surface density fluctuations. Thus, to analyze the structure and evolution of the hot spots in the disk, we will primarily focus on the surface density $\Sigma(r,\phi,t)$ defined as

$$\Sigma(r,\phi,t) = \int d\theta \sqrt{g_{\theta\theta}}\rho(r,\theta,\phi,t). \tag{31}$$

For each annulus of constant radius and time, we subtract out the mean and normalize the surface density such that

$$\int d\phi \bar{\Sigma}(r,\phi,t) = 0 \tag{32a}$$

and

$$\int d\phi \bar{\Sigma}^2(r,\phi,t) = 1. \tag{32b}$$

To estimate the distribution of hot spot lifetimes, we calculate the correlation between the surface density maps at two different times. The time-symmetric overlap function is defined as

$$c(r,t,\Delta t) \equiv \frac{1}{2}\left[\int d\phi \bar{\Sigma}(r,\phi,t)\bar{\Sigma}(r^+,\phi^+,t+\Delta t)\right]^{1/2} + \frac{1}{2}\left[\int d\phi \bar{\Sigma}(r,\phi,t)\bar{\Sigma}(r^-,\phi^-,t-\Delta t)\right]^{1/2} \tag{33}$$

so that $c(r,t,0) = 1$. Here the advanced and retarded coordinates (r^+,ϕ^+) and (r^-,ϕ^-) are the spatial positions where a clump of matter at (r,ϕ,t) would be at times $t+\Delta t$ or $t-\Delta t$. These coordinates are determined by integrating the time- and azimuth-averaged fluid velocities at each radius in the disk. For example,

$$r^+ = r + \int_0^{\Delta t} \bar{v}^r[r(t)]dt \tag{34a}$$

and

$$\phi^+ = \phi + \int_0^{\Delta t} \bar{v}^\phi[r(t)]dt, \tag{34b}$$

where $\bar{v}^i(r)$ is the average fluid transport velocity (dx^i/dt) in the disk midplane. This more detailed approach is necessary for the inner disk where the radial velocities are not negligible and the orbital velocities change rapidly with r.

Now the complete correlation function at each radius can be given by averaging over the duration of the simulation:

$$C(r, \Delta t) = \frac{1}{t_f - t_i} \int_{t_i}^{t_f} c(r, t, \Delta t)dt. \tag{35}$$

The correlation is normalized so that if there is no evolution of the disk perturbations in the local fluid frame, then $C(r, \Delta t) = 1$. Of course, in practice there *is* significant evolution, and therefore we expect $C(r, \Delta t \to \infty) \to 0$ as the turbulent disk at late times is completely uncorrelated with the initial conditions. Then the characteristic decay time for $C(r, \Delta t)$ is a measure of the lifetime of a typical hot spot at that radius.

To compare lifetimes at different radii, we rescale the time delay Δt by the orbital period ($2\pi/\Omega$) at each radius. These correlation functions are plotted in Figure 10 for a range of r, including the region inside the ISCO. Remarkably, we find very similar behaviors for hot spots over a large range of radii ($3M \leq r \leq 25M$), with nearly exponential decay at early times with a time constant of

$$T_{\rm life} \approx 0.3 T_{\rm orb}. \tag{36}$$

This behavior is quite similar to the simple hot spot model developed by Schnittman (2005) where hot spots were created and destroyed around a single radius with random phases and exponentially-distributed lifetimes, naturally giving Lorentzian peaks in the power spectrum at the geodesic frequencies corresponding to the chosen "resonant" radius. However, for the MHD simulations, there appears to be no special radius, and even for the potential QPOs discussed above in Section 4.2, the geodesic hot spot model may not be appropriate.

In addition to the typical lifetimes for the density perturbations, we have also analyzed the characteristic *spatial* dimensions of the hot spots in r and ϕ. We would ideally like to follow each hot spot in its own locally non-rotating frame to remove the shearing effects of differential rotation in the disk. This would give a better estimate for the "true" size and shape of the local perturbations without including the global dynamics of the disk. In fact, there is a simple analytic transformation that can be applied to remove the shear effects from the entire disk simultaneously (at least the region that behaves in a roughly Keplerian way; $6M \lesssim r \lesssim 25M$ in our simulation).

To construct this transformation, consider a sheared blob extending in r from $r_0 - dr$ to $r_0 + dr$. The orientation of the blob at its creation can be estimated by shifting each point backwards in time by $T_{\rm life}(r_0)$. This amounts to a shift in azimuth of $\Delta\phi(r_0 \pm dr) = -T_{\rm life}(r_0)\Omega(r_0 \pm dr)$ or

$$\frac{d}{dr}\Delta\phi(r) = -T_{\rm life}(r)\frac{d}{dr}\Omega(r) \approx 0.9\left(\frac{\pi}{r}\right), \tag{37}$$

where we have assumed a Keplerian orbital velocity of $\Omega(r) \sim r^{-3/2}$ (recall this relation holds even for relativistic orbits in the Schwarzschild metric). Integrating equation (37) gives the azimuthal coordinate transformation $\varphi(r) = \phi(r) - \Delta\phi$, with

$$\Delta\phi(r) = 0.9\pi \ln\frac{r}{r_{\rm in}}. \tag{38}$$

Figure 11 shows the effect of applying this transformation to a snapshot of the accretion disk surface density. In the "before" image (Fig. 11a), the differential rotation is clearly present, as material in the inner disk gets swept along faster, and shears the overdensities into diagonal arcs in the (r, ϕ) plane. After performing the transformation defined above, the blobs appear rather uniform in shape and orientation (Fig. 11b). Finally, to compare the aspect ratios of blobs at different radii in the disk, we perform a simple coordinate transformation in the radial direction as well:

$$\eta(r) \equiv \ln(r/r_{\rm in}). \tag{39}$$

As we will show, these coordinates highlight the invariance of $\delta r/r$ for density perturbations throughout the disk.

Given a surface density map $\Sigma(\eta, \varphi, t)$ for each frame in time of the simulation, we normalize to get $\bar{\Sigma}$ as in equation (32) for all values of η and t. Then, for each time step we create two partial Fourier transforms, one in each spatial direction:

$$\tilde{\Sigma}(k_\eta, \varphi, t) = \frac{1}{2\pi}\int d\eta \bar{\Sigma}(\eta, \varphi, t)e^{-ik_\eta \eta} \tag{40a}$$

and

$$\tilde{\Sigma}(\eta, k_\varphi, t) = \frac{1}{2\pi}\int d\varphi \bar{\Sigma}(\eta, \varphi, t)e^{-ik_\varphi \varphi}. \tag{40b}$$

Averaging (in quadrature) over the duration of the simulation gives

$$\tilde{\Sigma}^2(k_\eta, \varphi) = \frac{1}{N_t}\sum_{j=1}^{N_t}\tilde{\Sigma}^2(k_\eta, \varphi, t_j) \tag{41a}$$

and

$$\tilde{\Sigma}^2(\eta, k_\varphi) = \frac{1}{N_t} \sum_{j=1}^{N_t} \tilde{\Sigma}^2(\eta, k_\varphi, t_j). \tag{41b}$$

With this construction, $\tilde{\Sigma}^2(k_\eta, \varphi)$ is the time-averaged power spectrum of radial variations at a fixed azimuth φ. Similarly, $\tilde{\Sigma}^2(\eta, k_\varphi)$ is the average azimuthal power for a fixed radius η. Figure 12 shows examples of these power spectra for representative values of η and φ. It is clear from these plots that the variance $k\tilde{\Sigma}^2(k)$ can be well approximated by a broken power law for variations in both the radial and azimuthal directions. The "break wavenumber" gives the characteristic maximum size of a density perturbation. Since we find the peak amplitude of the perturbations increases with their spatial extent, it is reasonable that the variance peaks at this break wavenumber.

For each value of φ we fit the power spectrum $\tilde{\Sigma}^2(k_\eta, \varphi)$ to a broken power-law in k_η and find an average value for $k_\eta^{\rm break} = 5.5$. Applying a similar method to $\tilde{\Sigma}^2(\eta, k_\varphi)$ gives an average azimuthal break wavenumber of $k_\varphi^{\rm break} = 15$. These values correspond to wavelengths of $\lambda_\eta = 0.26$ (giving $\delta r/r = e^{\lambda_\eta} - 1 = 0.3$) and $\lambda_\varphi = 25°$, or an aspect ratio for the hot spots of $a/b = r\delta\varphi/dr \approx 1.5$. As we saw above in Section 4.2, the light curves from the top- and bottom-views of the optically thick disk are highly correlated. This suggests that the vertical extent of the density perturbations is limited by the disk thickness, which has a roughly constant value of $\delta z/r \approx 0.15$ (see Fig. 1d). Thus the density perturbations throughout the disk are self-similar ellipsoids with principle axes in the ratio $(r\delta\varphi : \delta r : \delta z) = (3 : 2 : 1)$.

To get a feeling for the stability of this wavenumber analysis, we plot in Figure 13 the best-fit parameters for the broken power-law at each slice in r or φ. Because of the rotational symmetry of the accretion system, it is expected that the fit parameters should be roughly constant in φ, so the curves in Figures 13d,e,f really give an estimate of the typical errors in our analysis. Somewhat more surprising is the near constancy of the fit parameters as a function of $\eta(r)$, emphasizing the fact that density perturbations at all radii have similar extent in azimuth. Lastly, the fact that the curves in Figure 13 do not appear as purely random fluctuations around the mean highlights the level of correlation between neighboring points in radius/azimuth.

6. DISCUSSION AND CONCLUSIONS

6.1. Summary of Findings

We have developed a generalized post-processor analysis tool that couples ray-tracing in the Kerr metric with global GRMHD simulations to produce light curves and power spectra

of accreting black holes. Using a variety of emission models, we probe different regions of the accretion disk and are better able to understand the underlying causes of time variability in the observed flux. The optically thick blackbody emission/absorption model is particularly useful for understanding the behavior of stellar-mass black holes in the Thermal Dominant state, which is characterized by a broad peak in the photon energy spectrum and very low levels of variability.

By fixing the black hole mass and accretion rate, we can convert from dimensionless code units to physical units for local fluid density. Then, assuming a radiation-pressure dominated gas in the inner disk, we can derive a physical temperature, and thus emissivity and absorption coefficients. The temperature and scale height of the disk compare well with the Novikov-Thorne predictions for a Schwarzschild black hole accreting at 50% $\dot{M}_{\rm Edd}$ with a small torque at the inner boundary. Despite the very different assumptions used in the analytic model and the computer simulations, the agreement is close enough to provide reasonable confidence in our conversion factors for density and temperature, and thus the conclusion that the simulations can be used to understand the Thermal Dominant state.

We have also developed new methods for analyzing the azimuthal structure of the accretion flow, determining the characteristic sizes and lifetimes of density perturbations. Over a large range of radii, we found the perturbations have a nearly exponential distribution of lifetimes, with $T_{\rm life} \approx 0.3 T_{\rm orb}$. Also, the characteristic shapes of the hot spots appears to be self-similar throughout the Keplerian regions of the disk, with $\delta\phi \approx 25^\circ$ and $\delta r/r \approx 0.3$.

From these short coherence times, it seems clear that the hot spots formed by MHD turbulence cannot survive long enough to produce the QPOs with quality factors of $Q \sim 5 - 10$ described in Schnittman (2005). Thus, if the transient high frequency QPOs with 2:3 frequency ratios identified in Section 4.2 are in fact robust signatures of GRMHD disk dynamics, they are most likely *not* produced by geodesic hot spots. This conclusion is supported by the fact that the QPOs are seen only from a single azimuthal viewing angle at any one time, perhaps suggesting some form of localized stationary wave in the disk. Finally, as part of our search for QPOs in the simulation data, we have also developed a general formalism for quantifying the significance of such features in simulations. These analytic results have been combined with Monte Carlo calculations to estimate our confidence limits for the QPOs at $\gtrsim 99.9\%$.

6.2. Comparison with Observations

As argued throughout this paper, we believe the MHD simulations most closely resemble the Thermal state of black hole X-ray binaries. This high-luminosity state is dominated by a broad thermal peak in the energy spectrum around $E \sim 0.7 - 1.5$ keV and a relatively small contribution from a steep power-law tail at higher energies (McClintock & Remillard 2005). The timing properties are characterized by a featureless power spectrum with $P(\nu) \sim \nu^{-\Gamma}$, with typical power-law index $\Gamma \approx 1$. The total power in this state is also small, with $(\mathrm{rms/mean})^2\ \mathrm{Hz}^{-1} \lesssim 10^{-3}$ above 1 Hz. Our simulated power spectra predict $\Gamma \approx 3 - 4$, which is rather greater than the typical slope. On the other hand, just this sort of steep power spectrum has been seen in the hard state observations of Belloni et al. (2006).

One robust prediction of the MHD/ray-tracing simulations is that, independent of the emission mechanism, the integrated rms power increases with disk inclination. In principle, this should be an observable prediction that could be tested with *RXTE* data from black hole binaries in the Thermal state. Of course, there are complications when comparing two different binary systems, including different relative contributions to the thermal peak or power-law tail, different total luminosities, and different black hole masses. However, with a sufficiently large number of observations, it should be possible to account for these variables and extract a correlation between integrated spectral power and inclination.

Other observational trends that might be explored with future MHD simulations include the observed linear relation between the X-ray flux and rms, which appears to extend to AGN as well (Uttley 2004; Uttley, McHardy, & Vaughan 2005). The timing properties of black hole binaries also vary greatly between their different spectral states, a general feature that is still not understood. To explore either of these relationships, we will require more sophisticated emission models and, quite likely, more simulation data with a broader range of disk parameters.

However, even with improved simulations and light curve models, there still exist some inherent difficulties in comparing theory with observation. *RXTE* typically requires thousands of seconds of observation to detect high-frequency QPOs at rms amplitudes of a few percent (limited largely by photon counting statistics and "dead time" corrections). To achieve a similar sensitivity, our MHD simulations require hundreds of hours corresponds to generate only 0.2 seconds of real time for a $10M_\odot$ black hole (we have perfect photon statistics, but are limited by the inherent variance in the power spectra of short time series). This is more closely analogous to the time scales of AGN observations. If we were to scale the simulation to an AGN of mass $10^8 M_\odot$, it would correspond to a sampling rate of every 40 minutes for roughly a month, assuming *perfect* photon counting statistics! Thus it should be no surprise that QPOs are still so difficult to detect with high significance from supermassive

black holes (Benlloch et al. 2001; Vaughan & Uttley 2005).

JDS acknowledges helpful discussions with Chris Reynolds and Tod Strohmayer. This work was supported by NSF grant PHY-0205155 and NASA grant NNG04GK77G (JFH), and by NSF grants AST-0205806 and AST-0313031 (JHK). The simulation described here was carried out on the DataStar system at SDSC.

REFERENCES

Abramowicz, M. A., & Kluzniak, W. 2001, A&A, 374, L19

Agol, E., & Krolik, J. H. 2000, ApJ, 528, 161

Armitage, P. J., & Reynolds, C. S. 2003, MNRAS, 341,1041

Balbus, S. A., & Hawley, J. F. 1998, Rev. Mod. Phys., 70, 1

Belanger, G., et al. 2006, ApJ, submitted; pre-print [astro-ph/0604337]

Belloni, T., et al. 2006, MNRAS, in press; pre-print [astro-ph/0603210]

Benlloch, S., et al. 2001, ApJ, 562, L121

De Villiers, J.-P., & Hawley, J. F. 2003, ApJ, 592, 1060

De Villiers, J.-P., Hawley, J. F., & Krolik, J. H. 2003, ApJ, 599, 1238

Done, C., Madejski, G. M., & Zycki, P. T. 2000, ApJ, 536, 213

Evans, C. R., & Hawley, J. F. 1988, ApJ, 332, 659

Fabian, A. C., et al. 1995, MNRAS, 277, L11

Fuerst, S. V., & Wu, K. 2004, A&A, 424, 733

Gammie, C. F., McKinney, J. C., & Toth, G. 2003, ApJ, 589, 444

Hawley, J. F., & Krolik, J. H. 2001, ApJ, 548, 348

Hawley, J. F., & Krolik, J. H. 2002, ApJ, 566, 164

Homan, J., et al. 2005, ApJ, 623, 383

King, R. E., et al. 2001, Phys. Rev. A, 64, 053810

Laor, A. 1991, ApJ, 376, 90

Machida, M., & Matsumoto, R. 2003, ApJ, 585, 429

Markowitz, A., & Edelson, R. 2004, ApJ, 617, 939

Matthews, A. P., & Scheuer, P. A. G. 1990, MNRAS, 242, 623

McClintock, J. E., & Remillard, R. A. 2005, in *Compact X-ray Sources*, ed. W. H. G. Lewin & M. van del Klis, preprint [astro-ph/0306213]

McKinney, J. C., & Gammie, C. F. 2004, ApJ, 611, 977

Miller, J. M., et al. 2001, ApJ, 563, 928

Miller, J. M., et al. 2002, ApJ, 578, 348

Petri, J. 2005, A&A, 439, 443

Pollak, G. D., Delameter, N. D., Nash, J. K., & Hammel, B. A. 1994, JQSRT, 51, 303

Rebusco, P. 2004, PASJ, 56, 553

Remillard, R. A., et al. 2002, ApJ, 580, 1030

Reynolds, C. S., & Begelman, M. C. 1997, ApJ, 488, 109

Rezzolla, L., Yoshida, S'i., Maccarone, T. J., & Zanotti, O. 2003, MSRAS, 344, L37

Rybicki, G. B., & Lightman, A. P. 1979, *Radiative Processes in Astrophysics* (New York: Wiley-Interscience)

Schnittman, J. D. 2005, ApJ, 621, 940

Schnittman, J. D., & Bertschinger, E. 2004, ApJ, 606, 1098

Schnittman, J. D., & Craxton, R. S. 2000, Phys. Plasmas 7, 2964

Schnittman, J. D., & Rezzolla, L. 2006, ApJ, 637, L113

Smith, M. D., Norman, M. L., Winkler, K.-H. A., & Smarr, L. 1985, MNRAS, 214, 67

Strohmayer, T. E. 2001a, ApJ, 552, L49

Strohmayer, T. E. 2001b, ApJ, 552, L169

Tregillis, I. L., Jones, T. W., & Ryu, D. 2001, ApJ, 577, 475

Uttley, P. 2004, MNRAS, 347, L61

Uttley, P., McHardy, I. M., & Vaughan, S. 2005, MNRAS, 359, 345

Vaughan, S., & Uttley, P. 2005, MNRAS, 362, 235

Vanden Berk, D. E., et al. 2001, AJ, 122, 549

van der Klis, M. 2005, in *Compact X-ray Sources*, ed. W. H. G. Lewin & M. van del Klis, preprint [astro-ph/0410551]

This preprint was prepared with the AAS LaTeX macros v5.2.

Table 1: RMS power for different emission models and inclination angles

Emission Model	Inclination (degrees)	Power > 10 Hz (% rms)	Power > 100 Hz (% rms)	$\Gamma(> 50$ Hz)
thin line	0	5.74	0.22	3.8
	45	6.96	0.36	4.5
	70	7.79	0.54	4.1
thick line	0	4.94	0.78	3.1
	45	4.51	1.33	3.2
	70	5.89	1.39	2.7
thermal	0	16.2	1.54	3.0
	45	13.6	2.22	3.3
	70	14.0	2.87	3.3

Table 2: RMS power from different annuli of the accretion disk with $i = 0°$ for the thermal emission model

radii (r/M)	Power > 10 Hz (% rms)	Power > 100 Hz (% rms)	$\Gamma(10 - 100$ Hz)	$\Gamma(100 - 1000$ Hz)
2-6	83.2	26.2	1.8	2.9
6-10	43.8	6.6	2.7	3.6
10-15	22.4	2.8	2.8	3.9
15-20	18.3	1.6	2.4	4.4
20-25	12.7	1.0	2.6	4.5

Fig. 1.— Density-weighted averages of hydrodynamic properties of the disk versus radius, for the MHD simulations (*solid*) and a steady-state Novikov-Thorne disk with a non-zero torque at the ISCO (*dashed*). The MHD disk behaves like an accretion disk only inside $r/M \approx 25$, where $\dot{M}(r)$ is roughly constant.

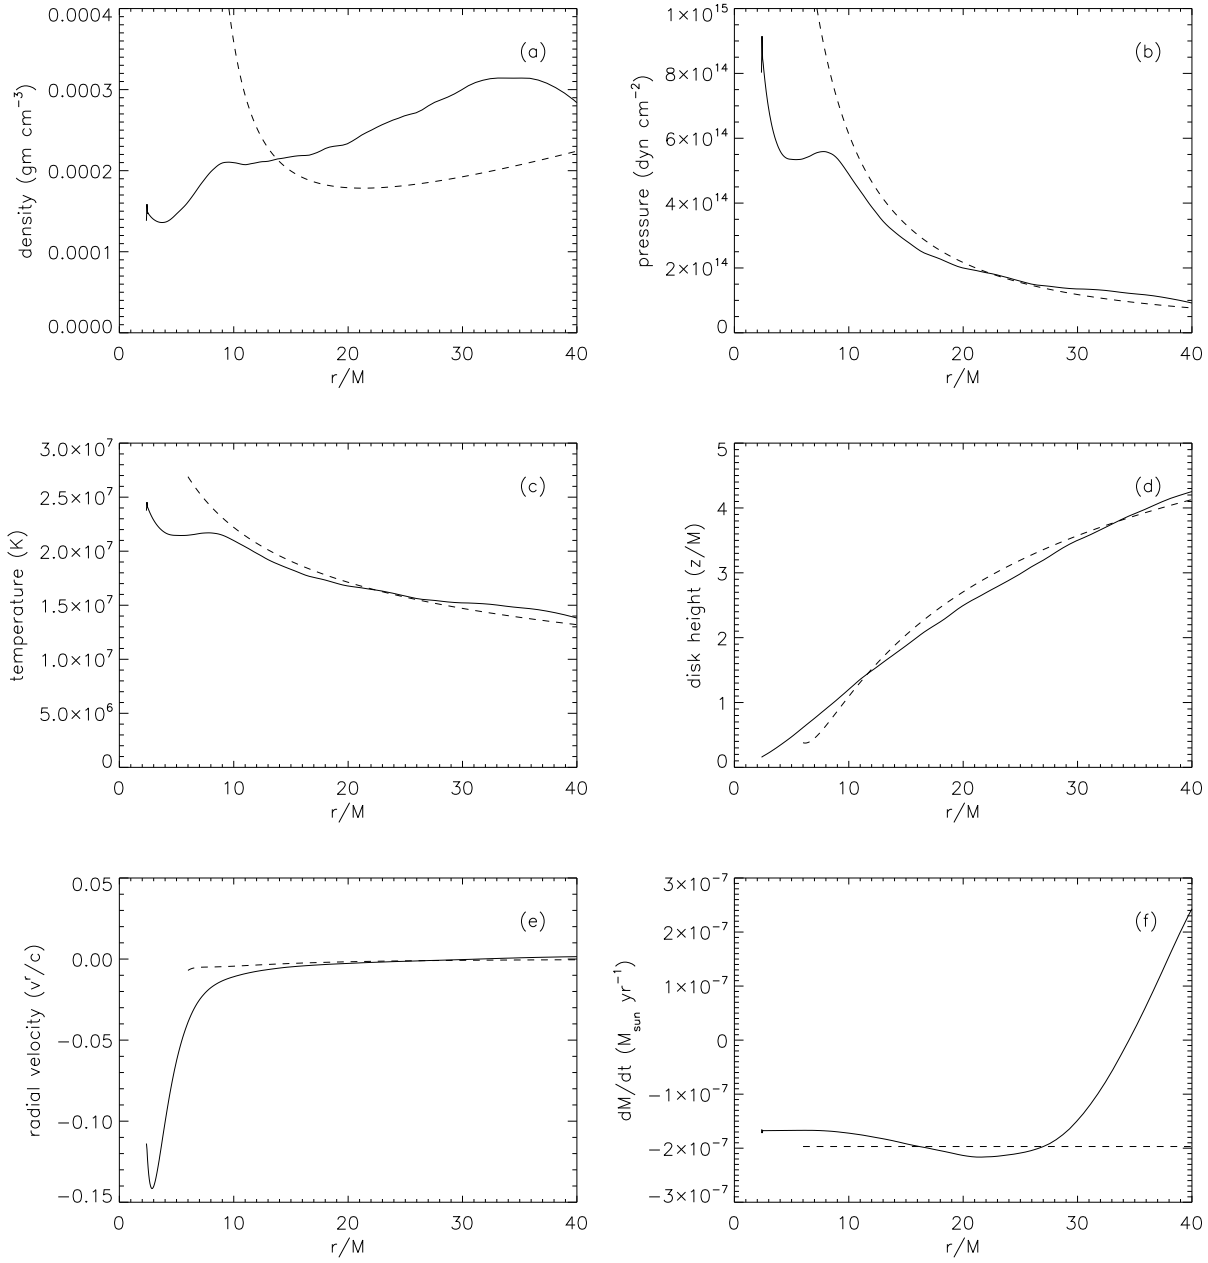

Fig. 2.— Schematic diagram of the ray-tracing procedure. Photon paths are integrated along geodesic trajectories from a distant observer to the region around the black hole. The positions and momenta are tabulated along the photon path and used to integrate the radiative transfer equation through the accretion disk (*shaded region*) towards the observer. Here the tabulation points are indicated schematically by diamonds; in practice much higher resolution is used. The photons either terminate at the black hole horizon (*thick circle*) or escape to infinity.

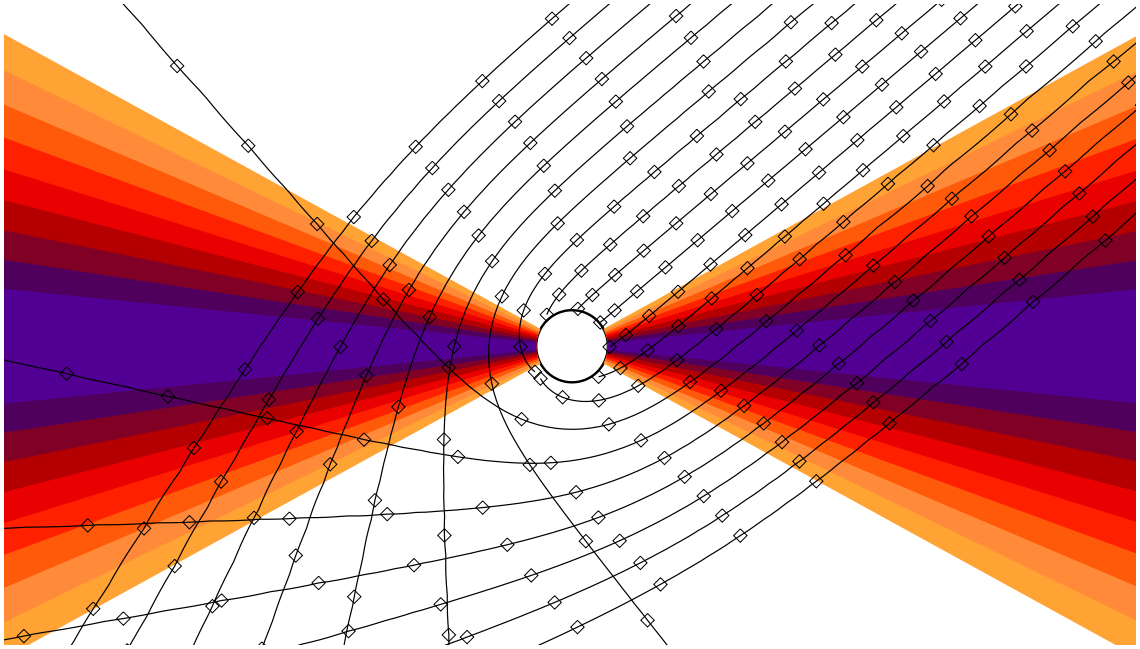

Fig. 3.— Snapshots of the inner regions of the MHD simulation, considering only emission inside of $r < 25M$. The three different emission models used are optically thin line (a, b), optically thick line (c, d), and thermal (e, f). The panels a, c, and e have viewer inclination angle of $i = 0°$ and the panels b, d, and e have $i = 70°$. For each panel, the color scheme represents a logarithmic scale, normalized to the peak intensity for that frame.

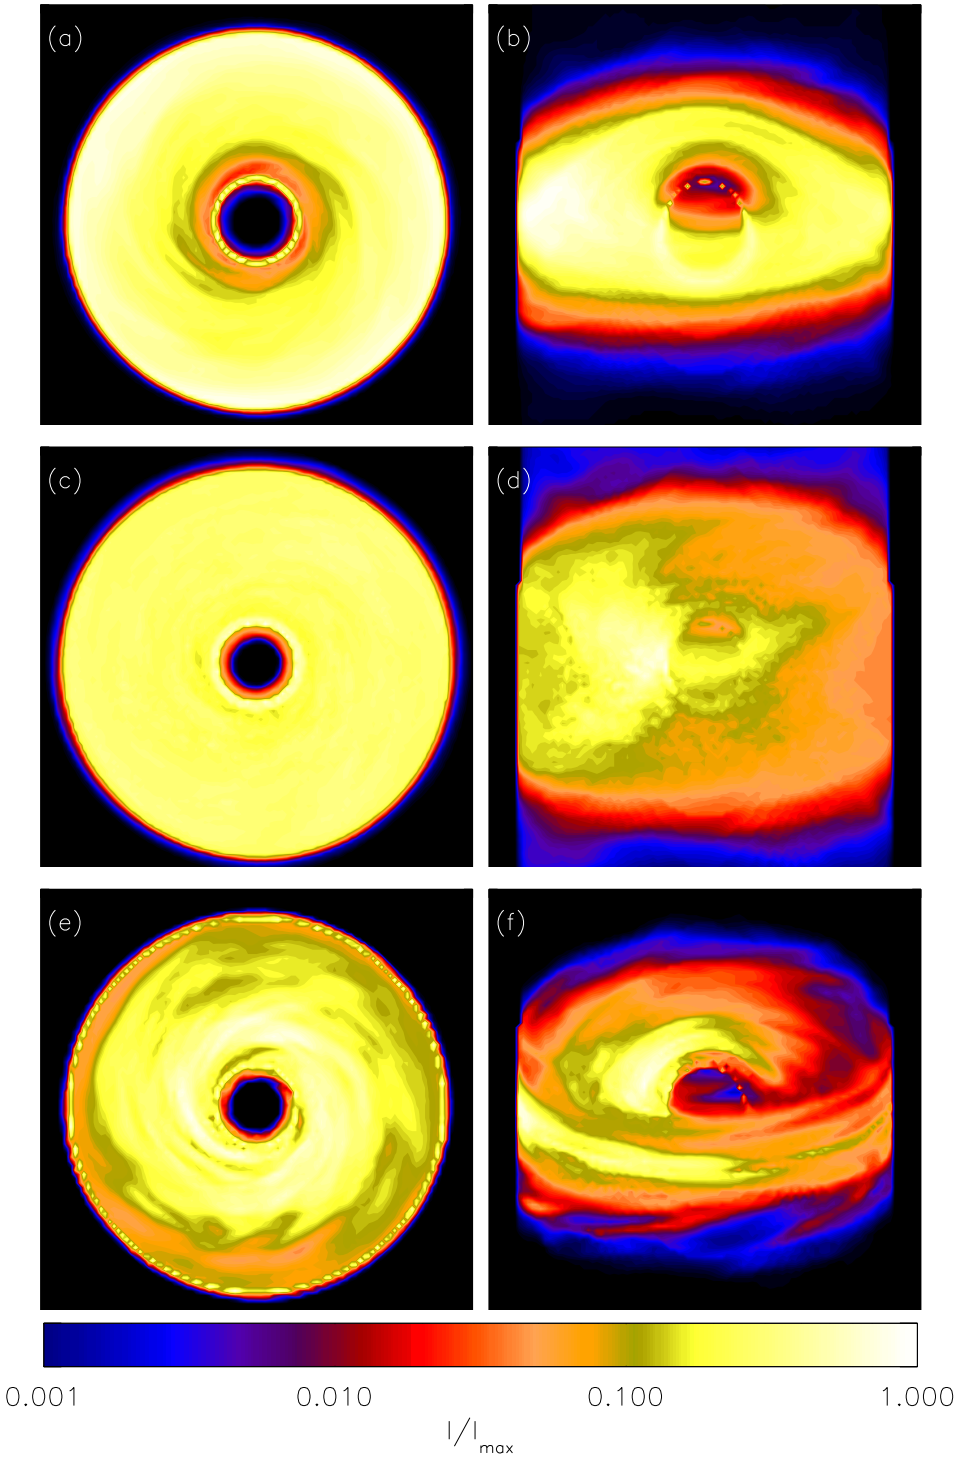

Fig. 4.— X-ray images of the inner regions of the MHD simulation, considering only emission inside of $r < 25M$. For a thermal emission model with $T_{\rm in} \sim 2 - 3$ keV, the three images are at photon energies of 1, 10, and 100 keV. The inner-most region, where the disk has the highest surface temperature, emits most of the low-energy X-rays (panel *a*). On the other hand, since the disk is optically thin to the highest-energy X-rays, which are produced in the regions of largest overdensity near the midplane, they can be seen equally well in the outer regions (panel *c*). In between (10 keV; panel *b*), the thermal peak produces the majority of the observed emission, and thus resembles the total intensity shown in Figure 3e. For each panel, the color scheme represents a logarithmic scale, normalized to the peak intensity for that frame.

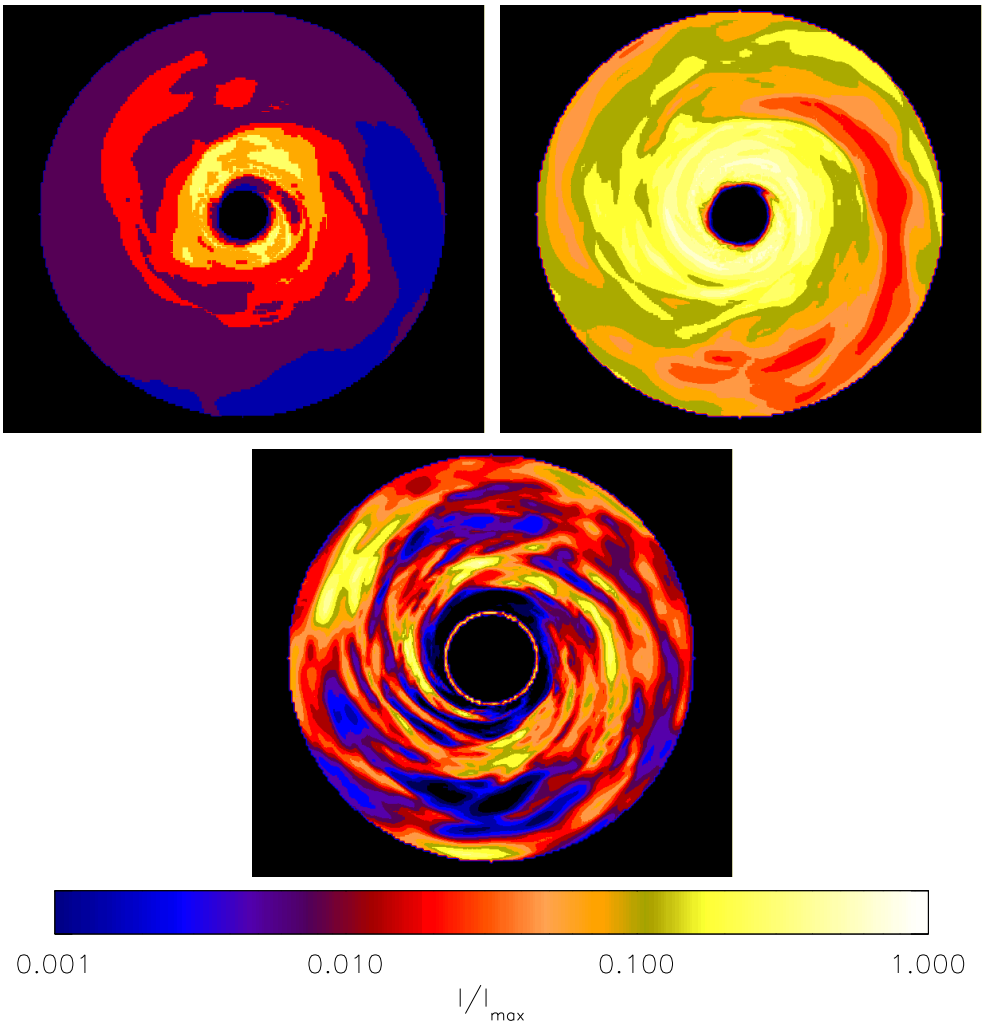

Fig. 5.— Integrated light curves for the three different emission models: (*a*) optically thin line, (*b*) optically thick line, and (*c*) thermal emission/absorption. The *solid* lines have inclination $i = 0°$, the *dashed* lines have $i = 45°$, and the *dot-dashed* lines have $i = 70°$. The light curves are all normalized to their mean amplitude, linear secular trends have been removed, and the higher inclination light curves have been shifted vertically in order to facilitate comparison. The time scale is set by assuming a black hole mass of $10M_{\odot}$.

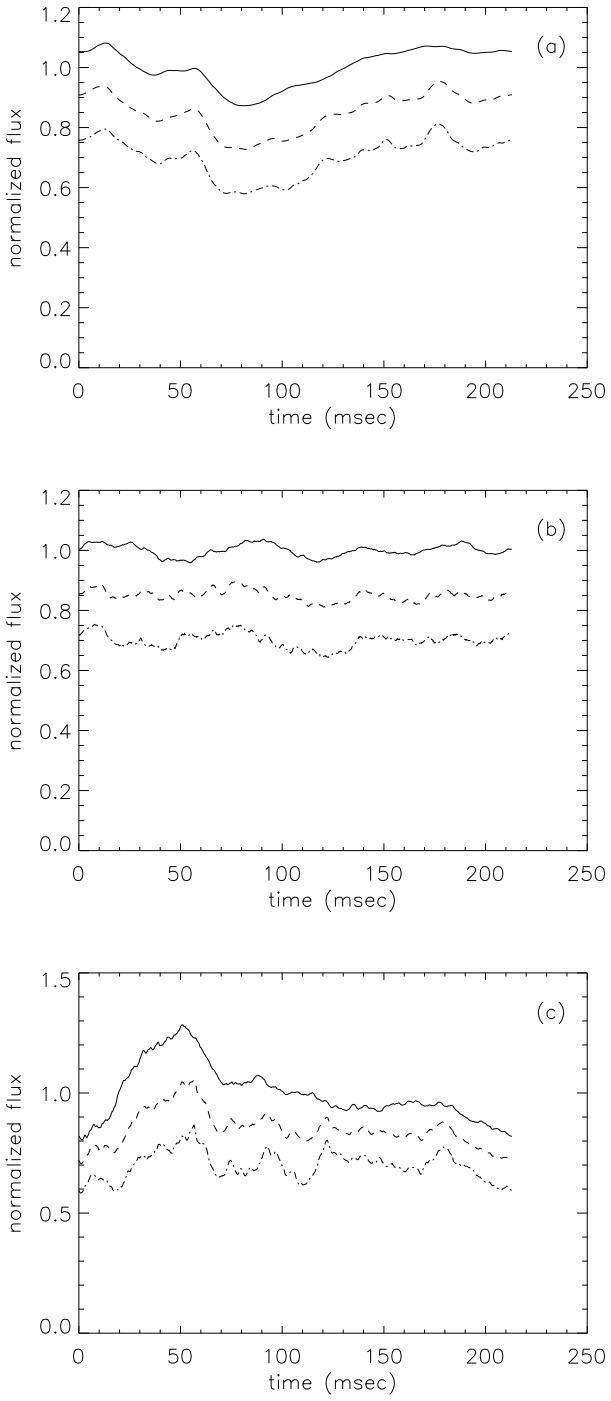

Fig. 6.— Power spectra for the light curves plotted in Figure 5. As in that Figure, the *solid* lines have inclination $i = 0°$, the *dashed* lines have $i = 45°$, and the *dot-dashed* lines have $i = 70°$.

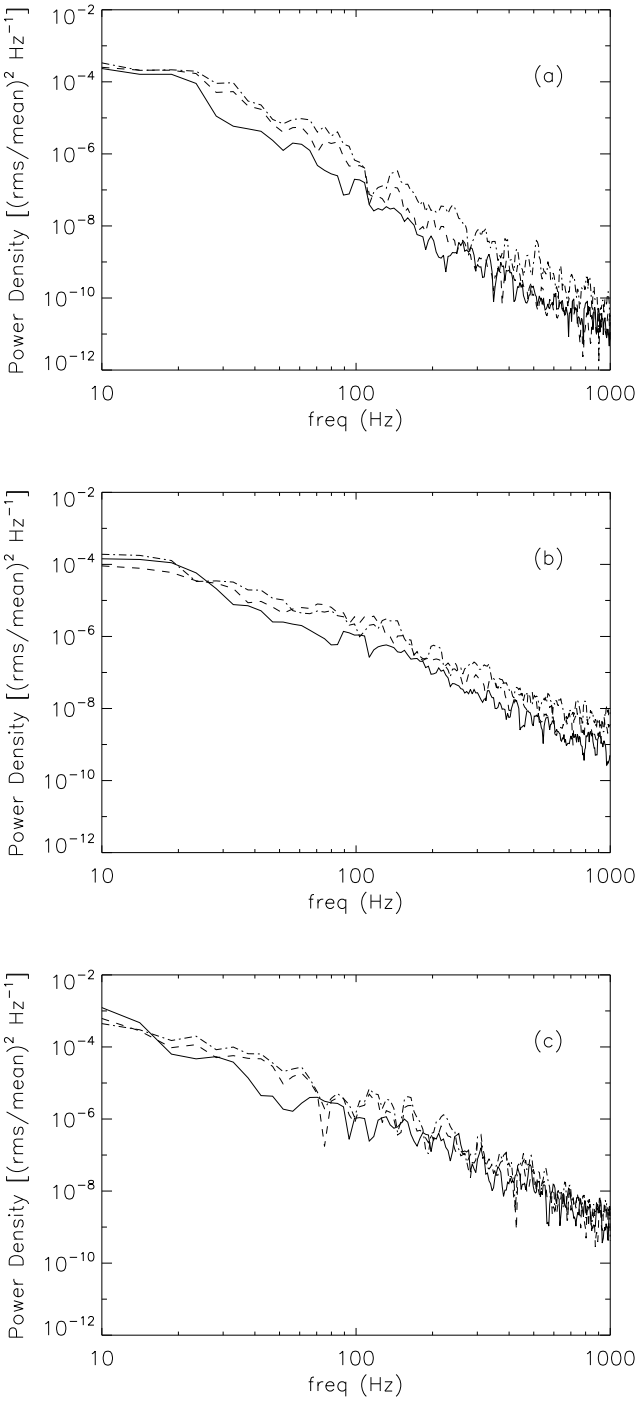

Fig. 7.— Power spectra for light curves produced by different annuli of the accretion disk, using the thermal emission model and an inclination of $i = 0°$. The inner-most regions of the disk produce larger amplitude fluctuations and contribute relatively more power at higher frequencies (flatter spectrum above ~ 100 Hz).

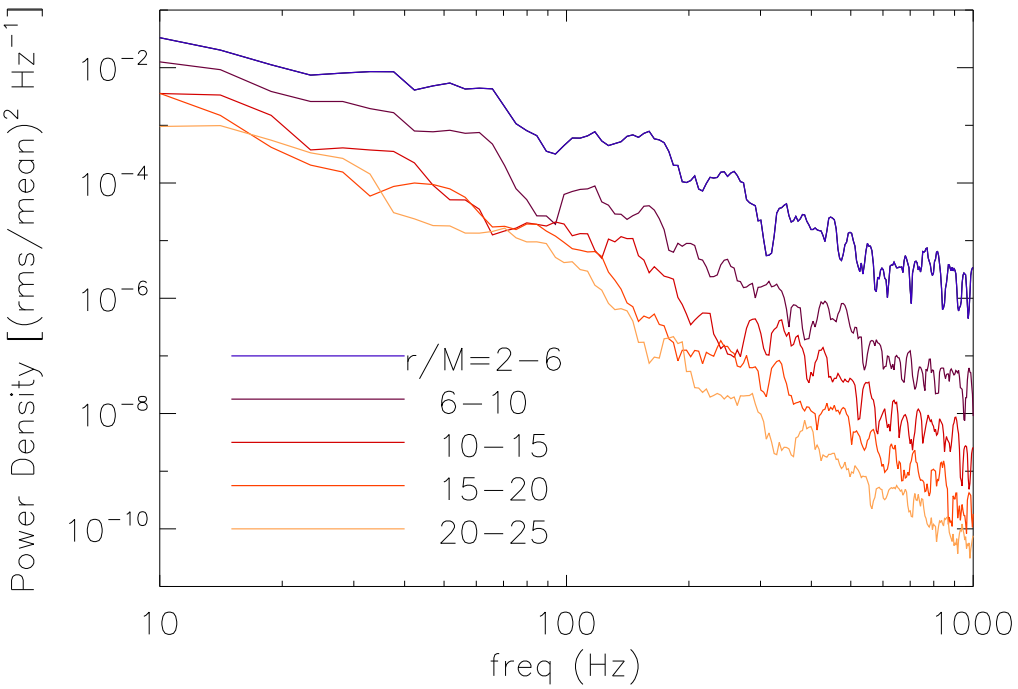

Fig. 8.— Excess power (in units of significance σ) above the "inherent" shape of the power spectrum, as defined by equation (30). The different plots correspond to different viewing angles of the same disk, all using the thermal emission model. The small ticks at the top of each plot correspond to integer multiples of a fundamental frequency ν_f: $2\nu_f$, $3\nu_f$, $5\nu_f$, and $6\nu_f$, where $3\nu_f = 230$ Hz, close to the geodesic orbital frequency at the ISCO (220 Hz).

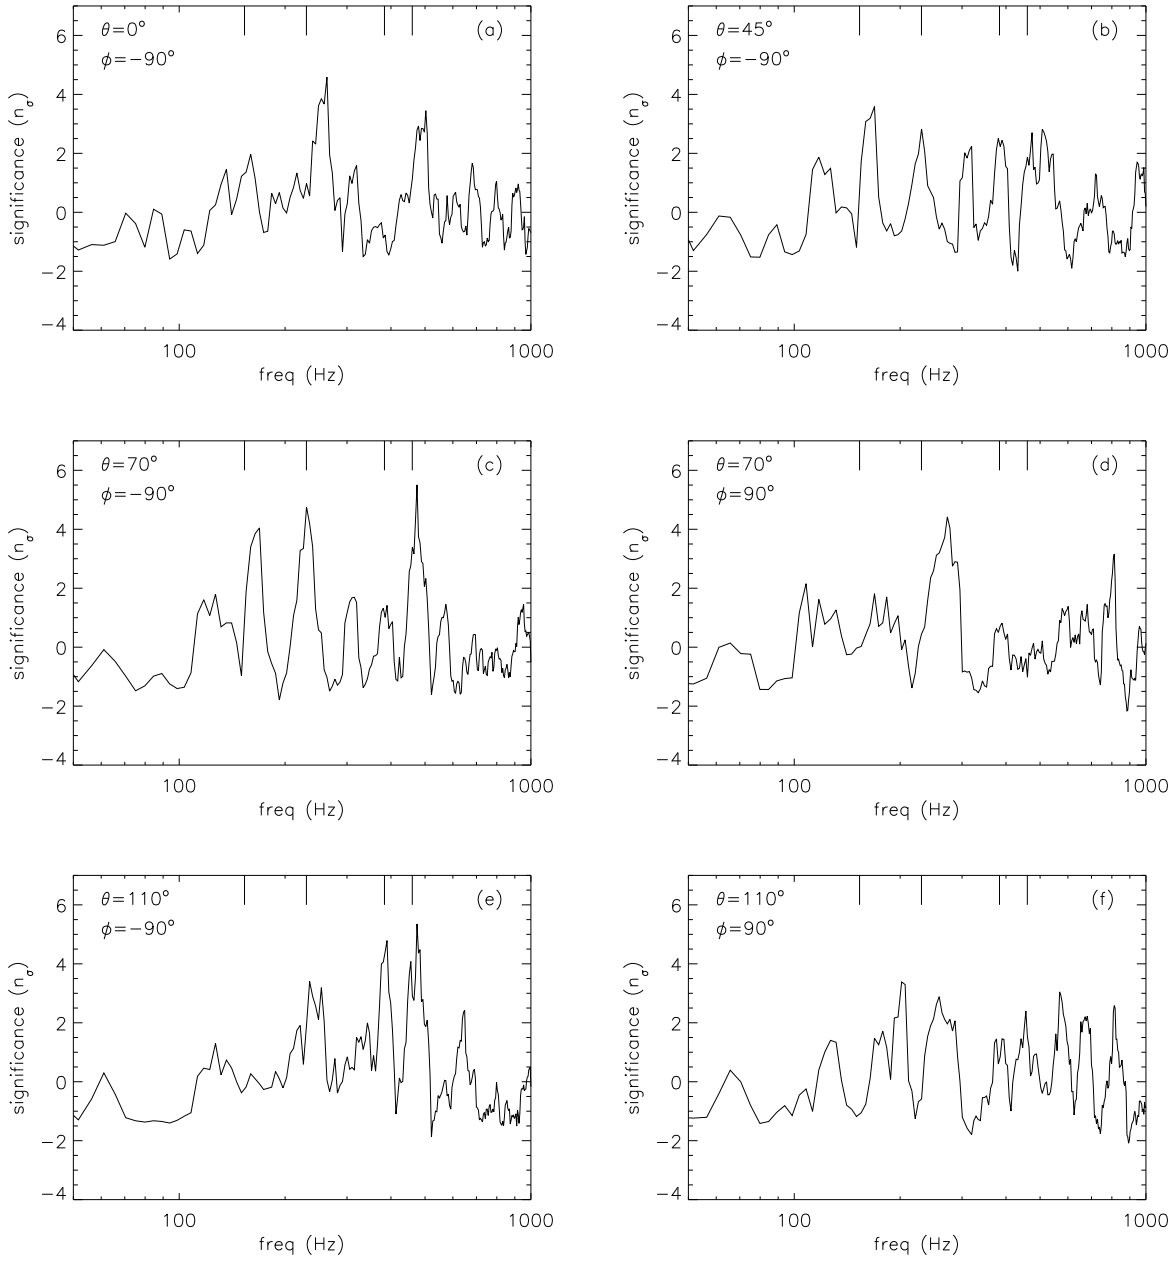

Fig. 9.— Probability of randomly producing a QPO at a given significance (n_σ), based on Monte Carlo calculations of a power-law power spectrum with fluctuations of individual frequency bins defined by an exponential probability distribution. As predicted by the analytic theory, we find 4σ results roughly $0.1 - 0.2\%$ of the time.

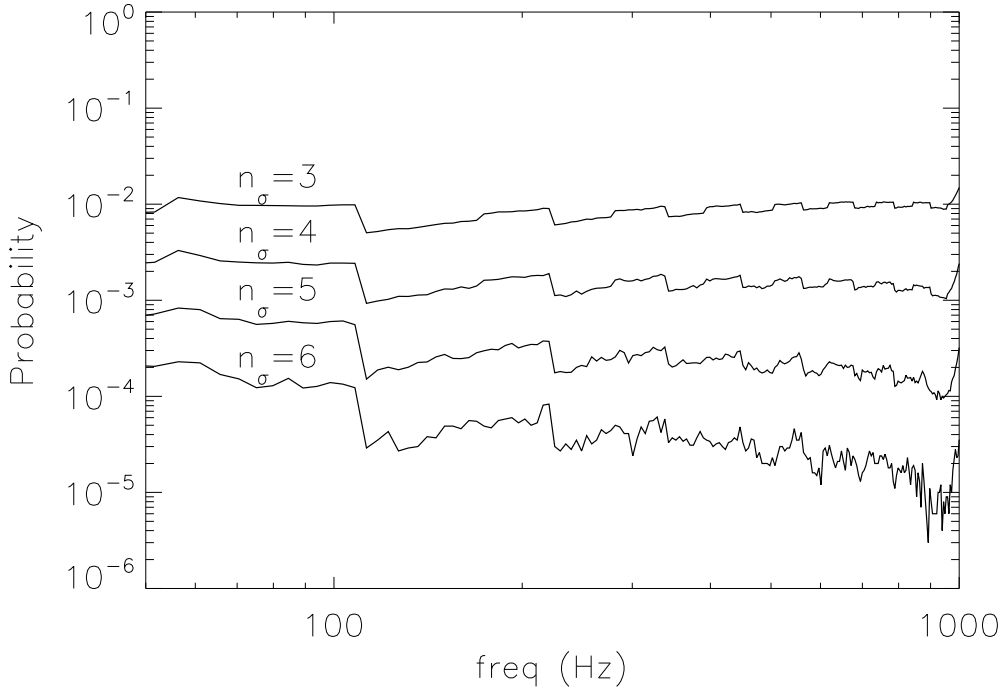

Fig. 10.— Cross-correlation of surface density fluctuations at different radii in the disk, as defined in equation (35). Outside of the ISCO, the typical density perturbation lasts a characteristic lifetime of about a quarter-orbit.

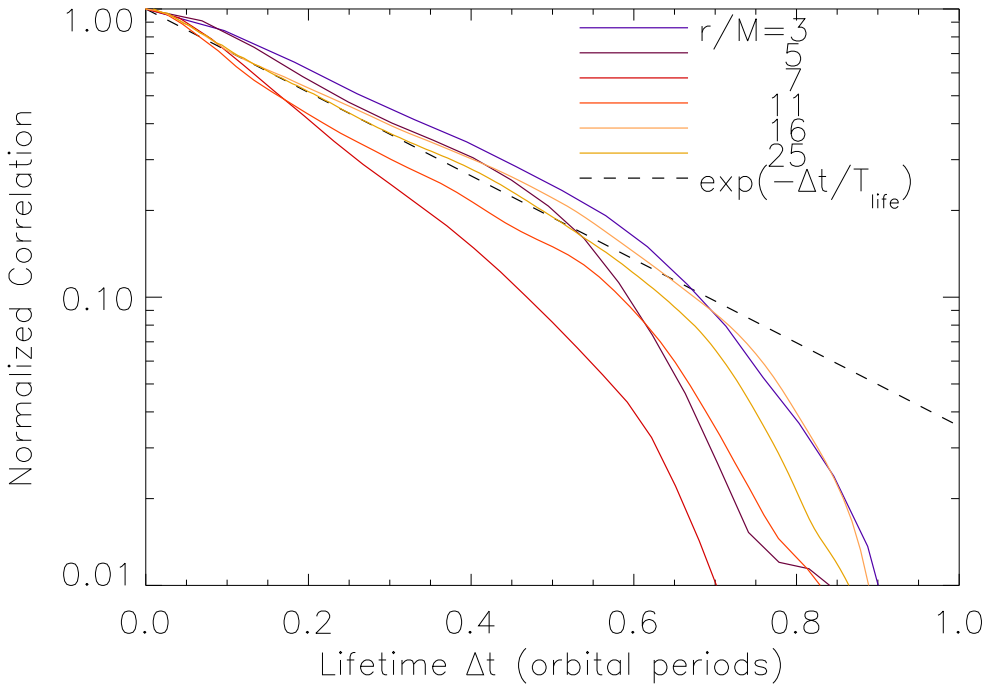

Fig. 11.— Surface density perturbations in the accretion disk at a single moment in time. In (a), clear shearing patterns are seen, caused by the differential rotation of the disk. In (b), we show the same perturbations, after "unshearing" the disk by translating each annulus of the disk by $\Delta\phi(r) = 0.9\pi \ln(r/r_{\rm in})$. The contours are spaced linearly and cover a range of $0.1 - 1.0$ times the maximum surface density.

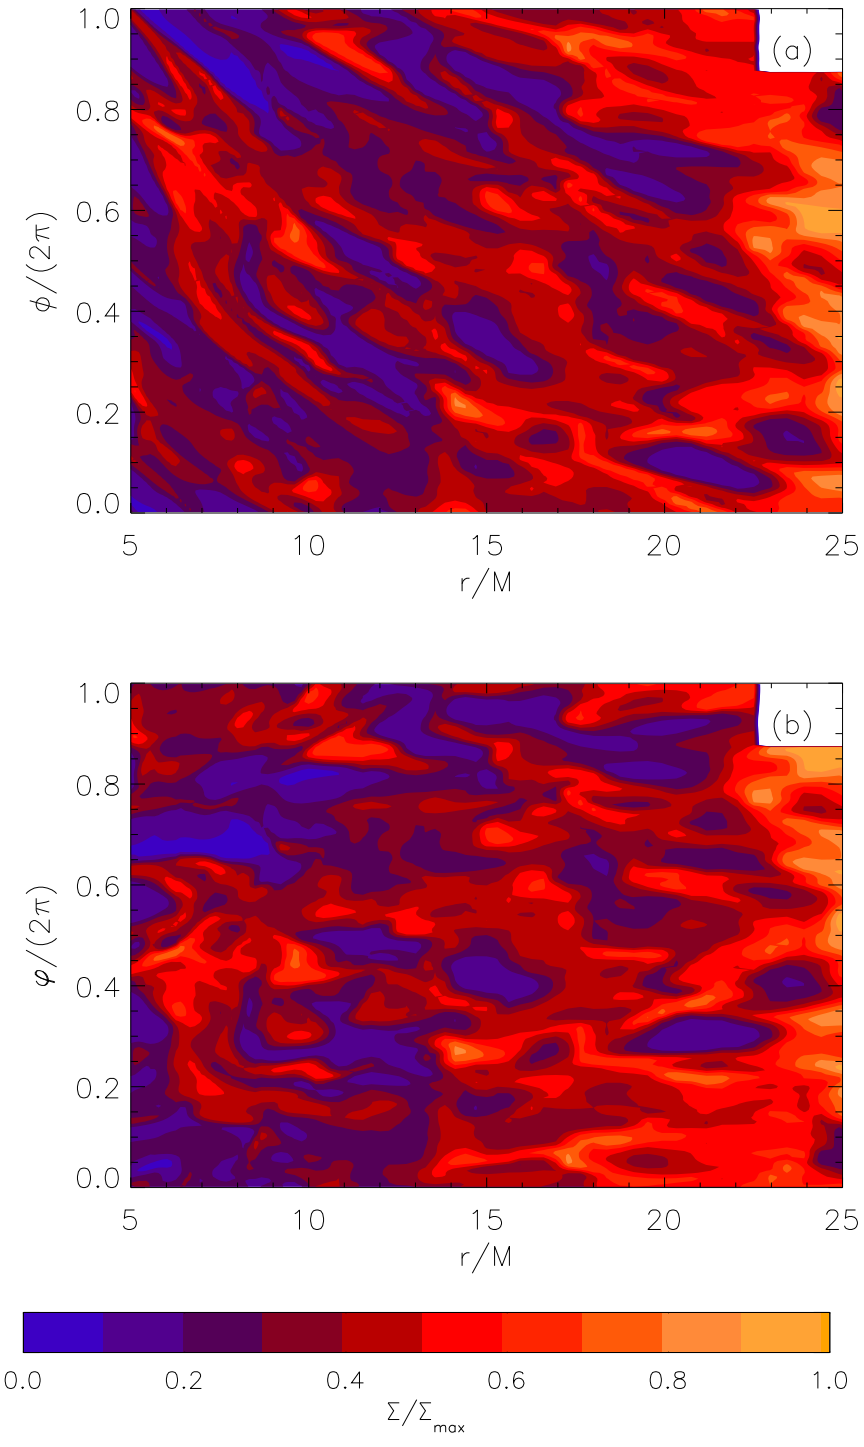

Fig. 12.— Variance in the normalized surface density as a function of wavenumber for (a) azimuthal modes and (b) radial modes. The characteristic size of a density perturbation is determined by the wavenumber of the power-law break.

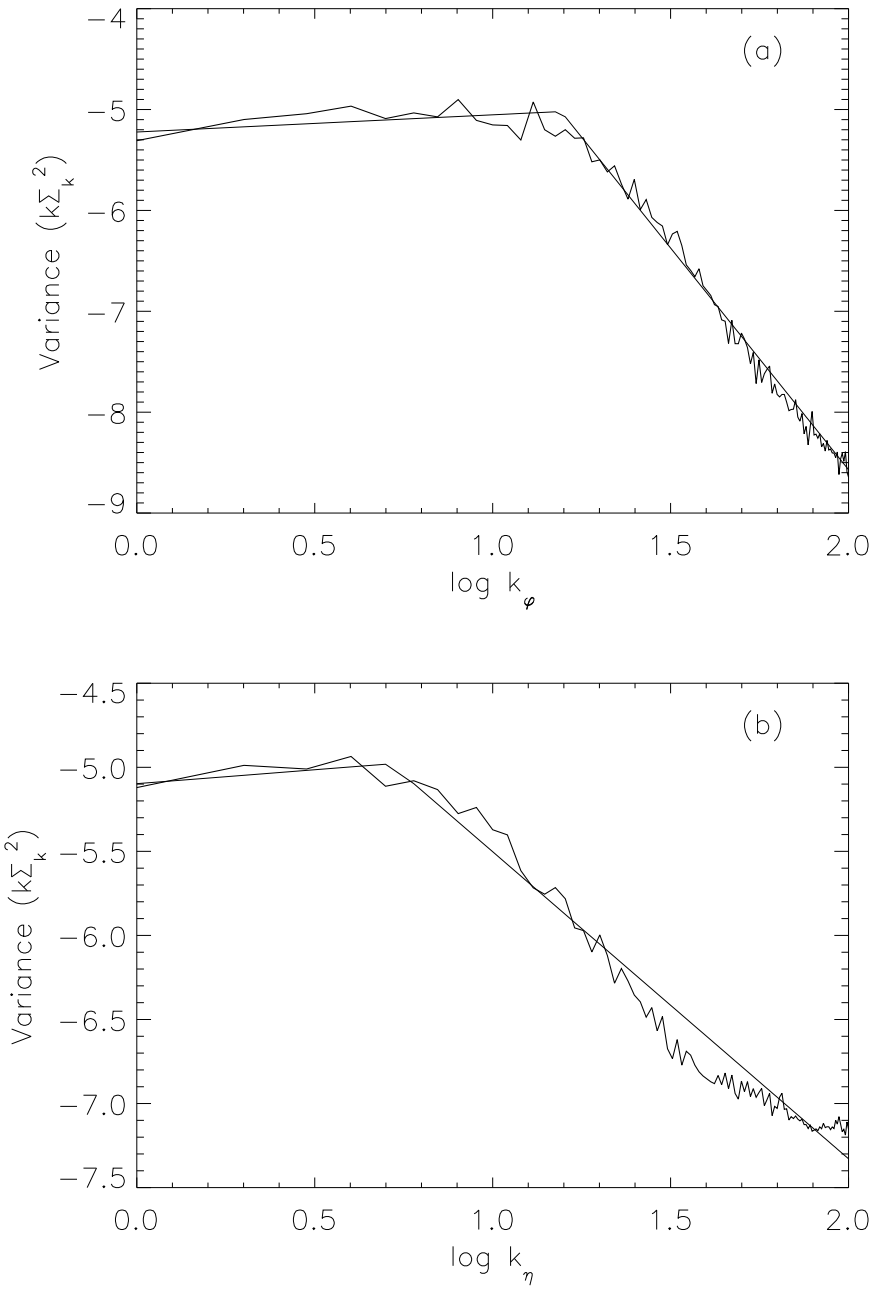

Fig. 13.— Best-fit parameters for the broken power-law curves plotted in Figure 12, where $\tilde{\Sigma}^2(k) \propto k^{-\Gamma_1}$ for $k < k^{\rm break}$ and $\tilde{\Sigma}^2(k) \propto k^{-\Gamma_2}$ for $k > k^{\rm break}$.